

Knowledge and Assumptions*

Brett Sherman and Gilbert Harman

August 11, 2010

Introduction

When epistemologists talk about knowledge, the discussions traditionally include only a small class of other epistemic notions: belief, justification, probability, truth. In this paper, we propose that epistemologists should include an additional epistemic notion into the mix, namely the notion of assuming or taking for granted.

Our starting point is pre-theoretical. We don't aim to show that current theories of knowledge are forced to take assumptions into account. Nor will we present a particular theory of knowledge. Our goal is to illustrate, in broad strokes, the epistemological picture you end up with when you start in a different place, one which we think is ordinary and straightforward.

*We have used material from Harman & Sherman (2004).

We begin by spelling out the ordinary and straightforward picture. After clarifying the notion of assumption that we appeal to, we go on to show how this picture can handle certain problems facing theories of knowledge that do not appeal to assumptions.

Taking Things for Granted

You often assume or take things for granted that you do not ordinarily take yourself to know to be the case. In thinking about where you will be next year, you may take for granted that you will not in the interim have died from a fatal heart attack or as a result of being hit by a bus. In thinking about what trips you will be able to afford to make, you may take for granted that a lottery ticket you have purchased is not the winning ticket. In thinking about whether you will be able to pick up a friend at the train station, you may take for granted that your car has not been stolen after you parked it outside the house a couple of hours ago. More generally you normally take for granted that you are not dreaming and are not a brain in a vat whose experience of a world of cars and trips and friends is artificially induced.

It is an interesting question whether you fully believe or accept things you take for granted in this way—that you will not have a heart attack in the coming year, that you will not be hit by a bus, that your ticket is not the winning ticket, that your car has not been stolen in the last hour, that you are not dreaming, that you are not a brain in a vat. You may very well fully

accept some of these things. In any event, you do not know them to be true.

But you may know other things that are in some sense supported by things you merely take for granted in this way. For example, you may know that your car is outside in front of your house. Your knowledge rests on various assumptions that you do not know but justifiably take for granted—that there is an external world including cars and houses, that you are not a brain in a vat who simply imagines you have a car and a house, and that no one has taken your car away since you parked it in front of your house an hour ago.

Part of the explanation of your knowing that your car is presently parked outside is that you justifiably (and truly) take it for granted that the car hasn't been stolen. That is not to say that the truth value of the knowledge claim is relative to the things you take for granted. The claim that you know your car is parked outside is not elliptical for anything more complex.

Of course, you cannot know something just because you are justified in taking it for granted. You are justified in taking it for granted that you are not dreaming, but that does not mean that, having taken that for granted, you know that you are not dreaming. Nor, having taken it for granted, are you then justified in believing you are not dreaming.

Knowledge that P which rests on justifiably taking it for granted that A is not just knowledge that, if A then P . You know that, if you are not a brain in a vat, you are not a brain in a vat. But you cannot come to know that

you are not a brain in a vat just because you are taking that assumption for granted.

Similarly, given your present situation, you may know that you will be in Paris next year. At the same time, you do not know you will not die before then. Your knowledge that you will be in Paris next year rests in part on the fact that you justifiably take for granted the assumption that you will not die before then. Alas, taking that for granted does not allow you to know you will not die before then.

This, it seems to us, is a natural way of describing the epistemic situation in the above cases. We won't present a theory of what it means to take something for granted. However, in the next section, we look at the role the notion of taking for granted plays in connection with inquiry in general, in order to help characterize the notion and distinguish it from related notions.

Inquiry, Investigation and Full Acceptance

Police open an investigation into a murder. They gather evidence and think about possible explanations of it that might help decide who did it. Once they are pretty confident who the murderer is *and confident that further investigation will not be worthwhile*, they end their inquiry.

In this case, their conclusion has a double aspect. Not only do they conclude *that* so-and-so committed the murder, they conclude *their investigation* into

the murder. They now treat this issue as settled.

In other instances police may close an investigation without having conclusively decided who the murderer is. They end inquiry because it seems to them no longer worthwhile to keep it open, given the costs of inquiry and their other priorities.

In the course of such an investigation police would not normally take seriously the possibility that so-and-so has a secret identical twin who actually committed the crime. Nor do they take seriously the possibility that they are “living in the matrix”. They simply assume they are not. Their inquiry rests on this and other background assumptions that they simply take for granted.

While the police are still investigating something, they are prepared to take seriously new evidence of a sort that they will not take seriously after the investigation is closed.

There are other examples of group inquiry aimed at answering a particular question—various sorts of scientific inquiry and various practical inquiries.

Similarly, a question arises for you as an individual. In a certain sense, you open your own investigation or inquiry (Dewey 1938). You think about it, perhaps gather evidence, and hopefully settle on your answer. That is, you come fully to accept a certain conclusion in a way that concludes your inquiry.

There can be full acceptance that is not the result of inquiry into the truth of what is accepted. You might fully believe or fully accept what you merely take for granted. Full acceptance—fully believing something—requires only that you treat the matter as settled.

One consequence of this is that fully believing something is not equivalent to accepting it as something you know.¹

You might tentatively accept something as a working hypothesis in order to see where it leads. In that case, you do not treat the matter as settled. If tentative acceptance of the hypothesis is sufficiently fruitful, you may become justified in fully accepting it.

Belief or full acceptance of P therefore involves two things. First, you allow yourself to use P as part of your starting point in further reasoning. Second, you take the issue to be settled in the sense that, when you fully accept P , you are no longer investigating whether P is true. (To be sure, you might continue investigating in order to get evidence that will convince others, for example, but you are no longer investigating in order to find out whether P is true.)

Reasoning tends to be conservative in the sense that you are normally justified in continuing fully to accept what you fully accept. Part of the explanation of this conservatism is that, once you take something to be settled, you need special reason to reopen inquiry into it before you can reconsider

¹Contrary to Harman 1986, p. 47.

it.

Of course, you may find yourself fully accepting P without having ever investigated whether P is true. For example, you realize that much of what you accept rests on such assumptions as that you are not a brain in a vat and you see no benefit in opening an investigation into whether such assumptions are true.

So, you might fully accept that you are not a brain in a vat without thinking that you know you are not a brain in a vat.

Strong Closure Principles

Intuitions about closure have persistently presented problems for theories of knowledge. One of the more interesting consequences of our picture of knowledge and assumptions is that we can explain some otherwise puzzling intuitions about closure straightforwardly.

It is not always true that, if you know that P and know that if P then Q , you should be able to know that Q . Q may be something you merely take for granted. You may know that your car is parked outside. You have that knowledge in part because you justifiably and truly take for granted that you are not dreaming. You also know that, if your car is parked outside, you are not merely dreaming that your car is parked outside. But you cannot come to know that you are not merely dreaming your car is parked outside

by taking it for granted that you are not dreaming.

A number of philosophers have defended “strong closure principles” for knowledge that conflict with the examples just given. Strong closure principles are incompatible with your knowing that your car is parked outside if you merely assume and do not know that the car has not been stolen since you parked it there.²

Philosophers have offered various motivations for strong closure principles. For example, such principles are said to be needed in order to account for how deduction can be a way of extending your knowledge (Williamson 2000, p. 117). Of course deduction is *sometimes* a way of extending your knowledge. But it is not *always* a way of extending your knowledge.

It is also a mistake to think that knowing something is relevantly like having a mathematical proof of it. If you have a mathematical proof that P and a mathematical proof that if P then Q , you can normally put these proofs together to give a mathematical proof that Q . Given a proof from assumption A to P and a proof from assumption B to “if P then Q ,” you can normally put these together to get a proof from assumptions A and B to Q . But the analogous principle for knowledge conflicts with the fact that you may know that your car is parked outside without knowing that you are not merely dreaming that it is.

²An example of a *weak* closure principle is one that holds not of knowledge but of what you either know or assume. Weak closure principles are compatible with the picture we present here, though we neither endorse nor reject them here.

There is a more basic worry about the apparent presupposition that deduction is a kind of inference, something you might do. Various philosophers assume, for example, that there is an activity of “competently deducing,” a kind of inferring that can provide one with knowledge (e.g., Hawthorne 2004a).

This is to confuse questions of implication with questions of inference. A deduction is a structured object, an abstract argument or proof. It is true that, in order to check or exhibit implications, you may sometimes construct arguments. And inference can be involved in that construction. But a deduction is the abstract argument constructed. Although constructing the argument is something you do, the deduction itself is not something you do. The deduction is not the constructing of the deduction.

Notice also that a given deduction can be constructed in various ways. You do not normally come up with a deductive argument by first thinking of premises, then thinking of intermediate steps in the relevant order, finally arriving at the conclusion. You often start with the conclusion and work backwards toward the premises. Or you start in the middle and work in both directions. Furthermore, although you might construct a deduction as part of a process of coming to accept its conclusion, you might also construct a deduction as part of a process of coming to accept one of its premises via an inference to the best explanation. The conclusion of a deduction is not in general the conclusion of an inference. The conclusion of an inference might be that a certain construction is indeed a valid deduction. The whole

argument is then the conclusion of the inference.

Inference and implication are conflated when principles of implication like Conjunction Introduction and Modus Ponens are treated as principles of inference that people might follow. This is a serious error. Principles of inference are normative and have a psychological subject matter. Principles of implication are not particularly normative and do not have a particularly psychological subject matter (Goldman 1986, Harman 1999).

Instead of saying, “deduction is a way of extending knowledge,” it is more accurate to say (A) inference is a way of extending knowledge and (B) inference can involve the construction of a deductive argument. But (A) and (B) are insufficient to support a strong closure principle for knowledge.

Inference and Implication

You may know that you are seeing a desk in part by taking for granted, but without knowing, that you are not a brain in a vat. This is a commonsensical way of describing an ordinary situation. You are normally justified in taking for granted that you are not a brain in a vat, but you do not know you are not a brain in a vat.

Suppose p is *that you are seeing a desk* and q is *that you are not a brain in a vat*. Concerning what we take to be the commonsensical view of this case Hawthorne (2004c) says,

The authors realize that a consequence of their proposal is that, very often, one knows p , does not know q , is in a position to competently infer q from p (without thereby losing knowledge that p), but is nevertheless in no position to know that q (510).

Clearly we do not “realize” any such thing, because we do not confuse recognizing an implication with inferring what’s implied. In the case under discussion you are *not* in a position “to competently infer” (the proposition that) q from (the proposition that) p , at least in the ordinary sense of “infer” in which to infer something is to come to believe it on the basis of that inference. No one who uses terms in this ordinary sense thinks you can competently *infer* that you are not a brain in a vat from your knowledge that you are seeing a desk, precisely because your knowledge rests on your taking for granted such things as that you are not a brain in a vat. Such an inference would clearly be *incompetent*, not competent.

A better formulation of the passage from Hawthorne might be: An obvious consequence of the authors’ proposal is that, very often, you know that p , do not know that q , know that the proposition that q is implied by the proposition that p , but are nevertheless in no position to know that q by recognizing this implication.”

In response to our complaint that strong closure principles confuse questions of implication with questions of inference, philosophers like Hawthorne sometimes say something like, “Part of the disagreement is surely terminological.

There is a perfectly good use of the English verb ‘deduce’ where it denotes an activity.”

Of course, there is indeed a perfectly good use of the English verb ‘deduce’ to refer to inference, as when Sherlock Holmes is said to deduce who the murderer is. But in this ordinary usage, to deduce a particular conclusion is to infer it. So competently to deduce something would be competently to infer it. If you cannot competently infer something, you cannot competently deduce it in this sense. There is a difference between inferring that q and merely recognizing that something you believe implies that q .

There is indeed a common philosophical use of ‘deduce’ which simply conflates these two distinct things. And we recognize that some philosophers use ‘deduce’ in that way (Hawthorne (2004abc), Field (2009), ...). But we do not agree that this is “a perfectly good use of the English verb ‘deduce’.” That is why we objected.

Attempts to Defend Strong Closure Principles

Philosophers defend strong closure principles against various attacks by arguing that denying such principles leads to highly undesirable consequences. Although we lack the space to address all of the arguments that have been offered, we will briefly indicate why we think these defenses all fail to take into account the way knowledge rests on assumptions.

Sometimes these defenses appeal to the claim that you can appropriately assert something only if you take yourself to know it. This particular claim obviously fails to hold in general. If the Nazis come to the door and ask where your brother is, you can appropriately assert something you know to be false. (See also discussion in Thomson, 2008, pp. 88-94.)

Furthermore, anyone who denies strong closure principles might very well assert something and refuse to assert a consequence of it, even when this consequence is pointed out. Hawthorne argues that this can lead to an odd conversation: Alice asserts that the animal in the cage is a zebra and agrees that, if the animal is a zebra, then it is not a cleverly disguised mule; however, she is not willing to agree that the animal is not a cleverly disguised mule.

But this is a mistake. Alice *accepts* that the animal is not a cleverly disguised mule—she assumes that. She just doesn't take herself to know it. So, we see no difficulty here.

Similarly, consider an "Equivalence Principle," according to which, if you know a priori that the propositions that P and that Q are equivalent and you know that P , then you are in a position to know that Q . Once it is acknowledged that knowledge can rest on assumptions, the Equivalence Principle has no more intuitive force than more general closure principles. Alice knows this animal is a zebra, on the assumption that it is not a cleverly disguised mule. And the animal's being a zebra is equivalent to its being a zebra and not a cleverly disguised mule. But, just as she cannot know on

the basis of her assumption that her assumption is correct, she is not in a position to know on the basis of that assumption that the animal is a zebra and not a cleverly disguised mule.

Another argument in defense of strong closure principles is that positive accounts of knowledge that deny strong closure, like Dretske (1970) and Nozick (1981), have significant problems. Conclusion: strong closure principles must be accepted. Such an argument is of course fallacious. Dretske and Nozick offer their theories as ways of accounting for certain intuitive knowledge claims. The fact that their particular theories have problems does not undermine the intuitive nature of those claims.

Furthermore, we don't know of any nontrivial positive analyses of knowledge without significant problems.

Variation in Acceptability of Knowledge Claims

Fred says he knows his car is parked out front. April challenges this by asking how he knows it hasn't been stolen since he came inside. Fred says it is very unlikely to have been stolen. April says this is not enough for him to know it hasn't been stolen. Fred agrees. Then what does he say?

He might say it is enough for him to take it for granted that his car hasn't been stolen and, assuming that what he is taking for granted is true, he knows his car is parked out front.

It is also possible that, in the face of April's challenge, Fred takes himself to be no longer justified in taking for granted that his car hasn't been stolen. And, Fred recognizes that, if he cannot take that for granted, he cannot know his car is parked out front.

Although the variation in the acceptability of Fred's knowledge claims stems specifically from the variation in what he takes for granted, the shiftiness exhibited in the face of April's challenge is an instance of a more general phenomenon.

In general, whenever a concept used in ordinary non-philosophical settings admits of some philosophical subtlety, it is easy for one to overlook the subtlety and be swayed by philosophical red herrings. For example, it is relatively easy to convince a college freshman that he or she doesn't act freely by pointing out to the freshman that all of his or her actions are caused by some prior state or event. It takes some work to see that freedom of the will needn't amount to freedom from causation. But when that subtle point isn't presented explicitly at the outset, it is easy to overlook it.

Whenever this kind of philosophical subtlety exists, one can pragmatically convey that a given factor is philosophically relevant simply by not suggesting otherwise. For example, by asking the freshman how his or her actions could be done freely given that the actions are caused, one pragmatically conveys that causation prevents freedom of the will.

The freshman will most likely follow the red herring. While we want to

explain the shift in the freshman's intuitions, we needn't explain them at face-value. That is, we needn't treat the freshman's intuitions (namely, that his or her actions aren't done freely) as accurate. We can explain the intuitions by explaining the mistake that the freshman makes.

Something similar happens in the case of knowledge. If knowledge can in fact rest on assumptions, it is nevertheless far from obvious that it can do so. So when April asks Fred how he could know where his car is parked given that he's assuming it hasn't been stolen, she pragmatically conveys that this assumption prevents Fred from knowing.

We think that Fred should stick to his guns. He should keep assuming that his car hasn't been stolen, and so he should keep taking himself to know that his car is parked outside. However, if he mistakenly follows the red herring, and stops taking for granted that his car hasn't been stolen, then we can explain why he ceases to take himself to know that his car is parked outside.

Inquiry is sometimes social. What you are justified in taking for granted in a conversation is a social matter. So, what you are justified in saying you know is sometimes a social matter.

This might help to explain variations in the acceptability of knowledge claims that some theorists have taken as evidence for a contextual element in knowledge claims.

Oddity

Claims like “I do not know whether or not the flight will crash but I do know that I will be at the arrivals gate following my flight” and “I do not know whether or not I am dreaming right now but I do know that I am in my kitchen” can sound odd.

But consider the following examples. The first two sound odd, but not the last two.

- (1) I do not know whether or not the flight will crash but I do know that I will be at the arrivals gate following my flight.
- (2) I do not know whether or not I am dreaming right now but I do know that I am in my kitchen.
- (3) In thinking about where I will be following my flight, I assume that things will go well and in particular that my plane will not crash. It seems to me that I am perfectly justified in taking that for granted. Given that assumption, I know that I will be at the arrivals gate following my flight.
- (4) I do not know whether or not I am dreaming. I think I am in the kitchen. In thinking this, I assume that I am not dreaming. As far as I can see I am justified in taking that for granted and, given that assumption, I know I am in the kitchen.

While (1) and (2) can seem odd as remarks out of the blue, (3) and (4) are

not at all odd. Furthermore (1) is not odd if preceded by (3) and (2) is not odd if preceded by (4). We suggest that whatever oddness there is to (1) and (2) is pragmatic.

We propose that the pragmatic principle responsible for the oddness is the same principle responsible for generating the philosophical red herrings exploited in the shifting knowledge attributions discussed above. According to that principle, for a concept whose nature is suitably subtle, one can pragmatically convey that a given factor is philosophically relevant for the application of that concept simply by not suggesting otherwise.

In the above cases, the non-obviousness of the relation between knowledge and assumptions is exploited in the odd assertions. It's odd to say, "I do not know whether or not the flight will crash, but I do know that I will be at the arrivals gate following my flight" because this pragmatically conveys, falsely on our view, that these are the only relevant epistemic states in the vicinity.

In other words, the suggestion from the odd assertion is that there is a connection between your knowing one thing and not knowing another, while raising the question how these can be so connected. There's a good answer to that question, but without articulating it, it's going to sound odd. That oddness results, in part, because the philosophy involved is non-obvious.

If this pragmatic explanation is correct, then we should be able to construct similarly odd conjunctions whenever a philosophical theory introduces pre-

cisifications to ordinary concepts. Consider the material eliminativist who denies that our ordinary concept of belief picks out a real mental state. Such a philosopher would endorse the truth of the following conjunction.

- (5) Folk Psychology is a false theory and I don't believe that Folk Psychology is a false theory.

Abominable! Of course, once we add in more of the details that explain the subtleties of the position, the result isn't odd at all.

- (6) Folk Psychology is a false theory. There are no such things as beliefs, so it is never strictly true to say that one believes something. For example, I don't *believe* that Folk Psychology is a false theory.

The fact that one can find consequences of a theory that sound odd when conjoined is not a compelling reason to abandon the theory.

Details

In our view, there are situations in which taking it for granted that S will lose the lottery can support knowing that S will not be able to afford an African safari this year but not support knowing that S will lose the lottery. But we have not provided specified general principles for determining when one can know that p because one takes it for granted that q but not know that r on that basis.

While we agree that it would be nice to provide such principles, we do not see that it is any objection to our commonsensical approach that we have failed to do so. (Is it an objection to our approach that we do not provide a definition of *know*?)

Conclusion

To sum up, we have argued that knowledge often rests on assumptions that you are justified in making even though you do not know those assumptions to be true.

References

- Dretske, F., (1970). "Epistemic Operators," *Journal of Philosophy* 67, pp. 1007-2023.
- Field, H. (2009): 'The Normative Role of Logic,' *Proceedings of the Aristotelian Society Supplementary Volume 83*, pp. 251-268.
- Goldman, A., (1967). "A causal theory of knowing," *Journal of Philosophy* **64**: 357-72.
- Goldman, A., (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Harman, G., (1976). "Practical reasoning," *Review of Metaphysics* **29**, pp. 431-463.

Harman, G., (1980). "Reasoning and evidence one does not possess," *Midwest Studies in Philosophy* **5**: 163-82.

Harman, G., (1999). "Rationality," in *Reasoning, Meaning, and Mind*. Oxford: Oxford University Press. Esp. pp. 18-23.

Harman, G., (2003). "Adler's evidentialism," unpublished presentation at the Pacific Division Meeting of the American Philosophical Association in San Francisco.

Harman, G. and Sherman, B., (2004). "Knowledge, Assumptions, Lotteries," *Philosophical Issues*, *14*, *Epistemology* (2004): 492-500.

Hawthorne, J., (2004a). *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hawthorne, J., (2004b). "Précis of *Knowledge and Lotteries*," *Philosophical Issues*, *14*, *Epistemology* (2004): 477-481.

Hawthorne, J., (2004c). "Replies," *Philosophical Issues*, *14*, *Epistemology* (2004): 510-523.

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

Skyrms, B., (1967). "The explication of 'X knows that p,'" *Journal of Philosophy* **64**: 373-89.

Thomson, J. J. (2008). *Normativity*. Chicago: Open Court.

Williamson, T., (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.