Philosophical Issues, 14, Epistemology, 2004

KNOWLEDGE, ASSUMPTIONS, LOTTERIES

Gilbert Harman and Brett Sherman
Princeton University

John Hawthorne's marvelous book contains a wealth of arguments and insights based on an impressive knowledge and understanding of contemporary discussion. We can address only a small aspect of the topic. In particular, we will offer our own answers to two questions about knowledge that he discusses.

(Q1) Why is it intuitively correct that one can know one will be in San Francisco next year but not intuitively correct that one can know one won't die in the interim?
(Q2) Why is it intuitively correct that one can know using statistical reasoning that a given coin is not heavily biased toward heads by tossing it a few times but not intuitively correct that one can know using statistical reasoning that one's lottery ticket is not the winning ticket?

The paper will appeal to two theses about knowledge that seem to us correct.

(T1) What one knows can and usually does rest on assumptions one justifiably takes for granted without knowing them to be true.
(T2) One knows only if one believes as one does because of something that settles the truth of that belief.

(T1) helps to answer (Q1) and to address issues about strong closure principles for knowledge. (T2) helps to answer (Q2).

Taking Things for Granted

We begin by observing that one often takes things for granted that one does not fully believe to be the case. For example, consider a perfectly

ordinary individual we will call "Sam." In thinking about where he will be next year, Sam takes for granted that he will not before then have suffered a fatal heart attack or have been killed by being hit by a bus. In thinking about what trips he will be able to afford to make, he takes for granted that the lottery ticket he has purchased is not the winning ticket. In thinking about whether he will be able to pick up his daughter at the train station, he takes for granted that his car has not been stolen after he parked it outside the house a couple of hours ago. More generally, he normally takes for granted that he is not dreaming and is not a brain in a vat whose experience of a world of cars and trips and daughters is artificially induced.

It seems to us that Sam does not fully believe any of the things he merely takes for granted: that he will not have a heart attack, that he will not be hit by a bus, that his ticket is not the winning ticket, that his car has not been stolen in the last hour, that he is not dreaming, that he is not a brain in a vat.

But Sam does believe other things that are in some sense supported by things he merely takes for granted. He believes his car is outside in front of his house, for example. Indeed, he *knows* that his car is parked outside in front of his house.

So (T1) seems highly intuitive to us.

(T1) What one knows can and usually does rest on assumptions one justifiably takes for granted without knowing them to be true.

Sam knows he owns a car that is presently outside in front of his house. It seems intuitively correct to us that Sam's knowledge rests on various assumptions that he does not know but justifiably takes for granted: that there is an external world including cars and houses, that he is not a brain in a vat who simply imagines he has a car and a house, and that no one has taken his car away since he parked it in front of his house an hour ago.

Part of the explanation of Sam's knowing that his car is presently parked outside is that he justifiably (and truly) takes 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 he takes for granted. His claim that he knows his car is parked outside is not elliptical for anything more complex.

Notice that Sam cannot know something just by taking it for granted. He is justified in taking it for granted that he is not dreaming, but that does not mean that, having taken that for granted, he knows that he is not dreaming. Nor, having taken it for granted, is he then justified in believing he is not dreaming.

This also means that knowledge that P which rests on justifiably taking it for granted that A is not just knowledge that, if A then P. Sam knows that, if he is not a brain in a vat, he is not a brain in a vat. But he is not

justified in coming to believe that he is not a brain in a vat and he cannot come to know that he is not a brain in a vat just because he is taking that assumption for granted.

Similarly, given Sam's present situation, intuitively he knows he will be in San Francisco next year. At the same time, it also seems intuitively correct to us that he does not know he will not die before then. Intuitively, his knowledge that he will be in San Francisco next year rests on the fact that he justifiably takes for granted the assumption that he will not die before then. Alas, the fact that he takes that for granted does not allow him to know he will not die before then.

Hawthorne (2004) notes Williamson's remark, "it is reasonable for me to believe that I shall not be run over by a bus tomorrow, even though I know that I do not know that I shall not be run over by a bus tomorrow" (Williamson 2000: 255). We do not agree with the first part of this. It seems to us that it is not reasonable for Sam fully to *believe* he won't be run over by a bus tomorrow. It is not something he can legitimately *infer* from its small likelihood. Nevertheless, it is reasonable for him to *take it for granted* that he won't be.

Hawthorne also remarks in a footnote, "In a situation where I have no clue what is going on, I may take certain things for granted in order to prevent paralysis, especially when I need to act quickly." He appears to think this is a special case, whereas it seems to us that one normally takes quite a lot of things for granted that one does not know to be the case. Hawthorne does not discuss how knowledge might depend on the fact that one has taken certain things for granted.

Strong Closure Principles

Intuitively, it is not generally the case that, if one knows that P and one knows that if P then Q, one should be able to know that Q. Q may be something one takes for granted on which one's knowledge that P depends. Sam knows that his car is parked outside. He has that knowledge in part because he justifiably and truly takes for granted that he is not dreaming. He also knows, that, if his car is parked outside, he is not merely dreaming that his car is parked outside. But, intuitively, he cannot come to know that he is not merely dreaming his car is parked outside by taking it for granted that he is not dreaming.

A number of theorists including Hawthorne defend closure principles for knowledge that conflict with the examples just given. We will refer to such principles as "strong closure principles." By a strong closure principle, we mean one that is incompatible with Sam's knowing that his car is parked outside if he merely assumes and cannot know that the car has not been stolen since he parked it there.

Hawthorne offers a motivation for such principles and then goes on to defend them against attacks. The motivation he offers is perfunctory.

> As Williamson remarks, such principles articulate what is an extremely intuitive idea, namely that "deduction is a way of extending one's knowledge" (2000, p. 117). (Consider, for example, the paradigmatic status of mathematical knowledge that proceeds by way of deductive truth.)

A footnote references a similar claim by Bertrand Russell. A little later, he quotes Feldman (1999, p. 95),

> To my mind, the idea that no version of the closure principle is true—that we can fail to know things that we knowingly deduce from other facts we know—is among the least plausible ideas to gain currency in epistemology in recent years.

Apart from appeals to authority, there are two arguments here.

First, there is Williamson's remark that "deduction is a way of extending one's knowledge." That is (sort of) correct, if it means that deduction is *sometimes* a way of extending one's knowledge. It is incorrect if it means that deduction is *always* a way of extending one's knowledge. But something like the second interpretation is what is needed to defend a strong closure principle. (Furthermore, the presupposition that deduction is something one does represents a serious philosophical error that we will say more about below.)

Second, there is the claim that knowing something is relevantly like having a mathematical proof of something. If you have a mathematical proof that P and a mathematical proof that if P then Q, then you can normally put these proofs together to give a mathematical proof that Q. And, given a proof from assumptions A to P and a proof from assumptions 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 such intuitions as that Sam knows that his car is parked outside without knowing that he is not merely dreaming that it is.

A more basic worry about the passage from Williamson is its presupposition that deduction is a kind of inference, something one does. Hawthorne apparently presupposes the same thing, since all the strong closure principles he takes seriously assume there is an activity of "competently deducing," a kind of inferring that can provide one with knowledge.

Surely, this confuses questions of implication with questions of inference. A deduction is a structured object, an abstract argument or proof. True, in order to check or exhibit implications, we sometimes construct arguments. And inference can be involved in that construction. But a deduction is the abstract argument that is constructed. Although constructing

the argument is something someone does, the deduction itself is not something someone does. The deduction is not the constructing of the deduction. Notice also that a given deduction can be constructed in various ways. One does 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. One often starts with the conclusion and works backwards toward the premises. Or one starts in the middle and works in both directions. Furthermore, although one might construct a deduction as part of a process of coming to accept its conclusion, one 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.)

The same conflation of inference and implication is involved when principles of implication like Conjunction Introduction and Modus Ponens are treated as principles of inference that people might follow. This is generally recognized to be a serious philosophical 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).

Where Williamson says, "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.

Hawthorne's Defense of Strong Closure Principles

In defending strong closure principles against various attacks, Hawthorne argues that denying such principles leads to highly undesirable consequences. Although we lack the space to address all of the arguments he offers, we will briefly indicate why we think his defenses fail to take into account the way knowledge rests on assumptions.

The first consequence that Hawthorne considers depends on the claim that one can appropriately assert something only if one takes oneself to know it. It follows from this that anyone who denies strong closure principles will assert certain things and will refuse to assert their consequences, even when those consequences are pointed out to them. 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. In fact, she assumes that. She just

doesn't take herself to know it and so does not assert it, although she can assert that it is something she accepts. So, we see no difficulty here.

A further consequence of denying strong closure principles, according to Hawthorne, is that one must give up at least one of three more restricted closure principles that he thinks would be intuitive even to one who denies the more general closure principles. One of these restricted principles is the 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.

However, 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.

Hawthorne's final argument in defense of strong closure principles is that the positive accounts of knowledge that have resulted from denying strong closure have significant problems of their own. We leave it to the reader to determine whether that applies to our intuitive claim that knowledge can rest on assumptions one does not know to be the case.

To summarize so far: (T1) is intuitively plausible, it explains (Q1), and it conflicts with strong closure principles. Strong closure principles are not intuitively plausible, they generate problems with respect to (Q1), and they have no good arguments in their defense.

Variation in Acceptability of Knowledge Claims

Sam 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. Sam says it is very unlikely to have been stolen. April says this is not enough for him to know it hasn't been stolen. Sam 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 more likely that, in the face of April's challenge, Sam takes himself to be no longer justified in taking for granted that his car hasn't been stolen. And Sam recognizes that, if he cannot take that for granted, he cannot know his car is parked out front.

So (T1) 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.

Here we agree with the account in Hawthorne's final chapter according to which the truth of a knowledge claim about a person can depend on what possibilities the person who is the subject of the knowledge claim ought to take seriously, because the possibilities that one ought to take seriously are those possibilities whose rejection one is not justified in taking for granted.

At this point we would like to say something more substantial about the important question of when one is justified in taking something for granted. But we will have to defer further discussion to another occasion.

Settling the Truth of a Belief

We now turn to the second of the questions (Q2) we want to address on this occasion.

(Q2) Why is it intuitively correct that one can know using statistical reasoning that a given coin is not heavily biased toward heads by tossing it a few times but not intuitively correct that one can know using statistical reasoning that one's lottery ticket is not the winning ticket?

Anscombe (1957, p. 56) plausibly suggests that the difference between intention and belief has to do with the "direction of fit." We suggest that one such difference between these types of mental state can be expressed in the following two principles.

Self-referential Guarantee in Belief Content: The content of a full belief that *P* is: "I am in this mental state because of something that settles it that *P*."
Self-referential Guarantee in Intention Content: The content of an intention that *P* is: "My being in this mental state settles it that *P*."

Harman (1976, 1980, 2003) defends these principles and their appeal to self-referential contents of mental states.

The first principle, concerning belief, suggests:

(T2) One knows only if one believes as one does because of something that settles the truth of that belief.

(T2) is a variant of the causal theory of knowledge. In its simplest form that theory says one knows that *P* only if one's belief that *P* is caused by the fact that *P*. This has to be modified to allow, e.g., for knowledge of the future: one knows that *P* only if one's belief that *P* is caused either by the

fact that *P* or by something that causes it to be the case that *P* (Goldman, 1967). A further modification is needed to allow for Skyrms' (1967) example in which Alice shoots Bob to death. A short time later Carol cuts off Bob's head. Then Dan comes along and sees that Bob is dead. Dan does not believe that Bob is dead because of something that caused Bob to be dead (since Carol's cutting off Bob's head is not the cause of death) but Dan does believe this because of something that settles it that Bob is dead.

Now, suppose one has two coins, a nickel and a quarter, one biased towards heads to the extent that the probability of its coming up heads when tossed is 0.9, the other unbiased, so that the probability of its coming up heads when tossed is 0.5. After tossing each coin fifty times and noting the outcomes, one comes to believe that the nickel is the unbiased coin, because that explains why about half the tosses of the coin came up heads. Suppose one is right. The nickel is the unbiased coin. One believes that the nickel is unbiased because of the statistical results. And those results are as they are because the nickel is unbiased—that is, because of something that settles the truth of one's belief. So, one believes as one does because of something that settles the truth of one's belief. One can therefore know in this way that the nickel is unbiased, according to (T2).

On the other hand, one cannot in the same way know that one's lottery ticket is not the winner. As in the coin example, one believes as one does for probabilistic reasons. But in this case, the statistical facts are not explained by what settles the truth of one's belief. Suppose that one's ticket is not the winner. The fact that the ticket is not the winner does not explain the fact that it was highly unlikely to be a winner. So, if one were to believe it is not the winner for purely probabilistic reasons, one would not believe as one does because of something that settles the truth of that belief. And so one's belief would not be knowledge according to (T2).

It seems to us that one cannot be justified in fully believing something unless one is justified in taking oneself to know it. So, one would not be justified in believing the ticket is not the winner. Nevertheless, one might be justified in taking that for granted when one considers whether one will be able to afford to go on an African Safari next year. Given the fact that one justifiably and truly takes that assumption for granted, one might know that one will not be able to afford to go. But that does not mean one can know one's ticket will lose.

Conclusion

In summary, we have put forward two claims, (T1) and (T2), in order to answer two questions, (Q1) and (Q2). (T1) helps to answer (Q1) and explains one thing wrong with strong closure principles for knowledge. (T2) helps to answer (Q2).[1]

Notes

1. We are indebted to Elizabeth Harman and John Hawthorne for useful comments on earlier drafts.

References

Anscombe, G. E. M., (1957). *Intention*. Oxford: Blackwell.
Carnap, R., (1950). "Empiricism, semantics, and ontology," *Revue internationale de philosophiie* **4**: 20–40.
Feldman, R., (1999). "Contextualism and skepticism," *Philosophical Perspectives* **13**: 91–114.
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**: 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.
Hawthorne, J., (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
Skyrms, B., (1967). "The explication of 'X knows that p,'" *Journal of Philosophy* **64**: 373–89.
Williamson, T., (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.