

Knowledge and Practical Interests

Jason Stanley

OXFORD

Introduction

A central part of epistemology, as traditionally conceived, consists of the study of the factors in virtue of which someone's true belief is an instance of knowledge. The factors that have been proposed in epistemology are typically ones that are *truth-conducive*, in the sense that their existence makes the belief more likely to be true, either objectively or from the point of view of the subject. Much of epistemology has been devoted to debates between advocates of differing truth-conducive factors. For example, epistemic internalists have argued that the additional truth-conducive factors are other beliefs. Epistemic externalists have argued that the relevant truth-conducive factors include the fact that the belief is the product of a reliable belief-forming mechanism. All of these debates are between theorists who hold that only truth-conducive factors are relevant to the question of what makes it the case that someone's true belief is an instance of knowledge.

It is no surprise that epistemologists have widely shared the assumption that the additional factors that make a true belief into knowledge are uniformly truth conducive (either objectively or from the point of view of the subject). The differences between true belief and knowledge are matters that fall within the purview of theoretical rationality, which many philosophers hold

to be guided solely by the normative purpose of discovering the truth. My purpose in this book is to challenge this conception of knowledge. I will argue that the factors that make true belief into knowledge include elements from practical rationality. One consequence of my arguments is that the distinction between practical and theoretical rationality is less clear than one might wish.

Someone's practical investment in the truth or falsity of her belief is completely irrelevant to truth conduciveness in any sense. From the traditional perspective, then, when someone has a true belief, whether that belief is genuine knowledge is independent of *the costs of being wrong*. My aim is to provide a systematic case against this thesis. I join several recent authors in arguing that our *practical interests* have epistemic significance.¹ There are cases in which two people are similarly situated, but one has knowledge, whereas the other does not, because one has greater practical investment in the truth or falsity of her beliefs. What makes true belief into knowledge is not entirely an epistemic matter.

This conclusion is bound to sound somewhat paradoxical, because there are two senses in which epistemologists are prone to use the term 'epistemic'. On one use of 'epistemic', it denotes truth-conducive factors, in the broad sense in which I have sketched above. On the other understanding of 'epistemic', it has to do with factors relevant to whether a true belief is knowledge. The thesis of this book is that, contrary to epistemological orthodoxy, these two usages of the term do not coincide. Using 'epistemic' in the first of these ways, then, the thesis of the book is that what makes true belief into knowledge is not entirely an epistemic matter.

The book is short, because many of the elements of my argument have already been set in place by those with different goals. In particular, *contextualists* about knowledge ascriptions have discovered many of the examples that suggest that whether

¹ See Fantl and McGrath (2002) and Hawthorne (2004: ch. 4).

a true belief is knowledge depends not just upon truth-conducive features of a situation, but on what is practically at stake.² However, contextualists generally share the widely held assumption that knowledge is not a matter of practical interests. So they have used these examples, together with the assumption, to argue for the thesis that a predicate such as 'knows that penguins waddle' denotes different knowledge properties on different occasions of use. Each of the resulting semantic contents is a property, possession of which does not depend upon practical interests. But which such property is denoted by a knowledge-attributing predicate depends upon practical factors, such as how much is at stake. In this way, the contextualist can explain the examples without violating the commonly shared assumption that knowledge is not a matter of practical interests.

Contextualists have generally been interested in establishing the context-sensitivity of knowledge ascriptions in order to use the insight in the resolution of various traditional philosophical problems, such as explaining away the persuasive force of skeptical arguments. They have tended not to consider explicitly the assumption that what makes true belief into knowledge is purely a matter of truth-conducive factors, in the sense described above. But the interest of the examples they have employed to argue for the context-sensitivity of knowledge ascriptions is precisely that, when taken at face value, they do suggest the falsity of this assumption. Once we see that knowledge ascriptions are not context-sensitive in any distinctively epistemological way, we are led by such examples to reject the common assumption that knowledge (to put it tendentiously) is a purely epistemic notion.

Here are the examples I will focus upon; they have largely been made famous by others.

Low Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way

² In particular, most of the examples have been discovered by Stewart Cohen and Keith DeRose.

home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning.'

High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, 'I guess you're right. I don't know that the bank will be open tomorrow.'

Low Attributor–High Subject Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Two weeks earlier, on a Saturday, Hannah went to the bank, where Jill saw her. Sarah points out to Hannah that banks do change their hours. Hannah utters, 'That's a good point. I guess I don't really know that the bank will be open on Saturday.' Coincidentally, Jill is thinking of going to the bank on Saturday, just for fun, to see if she meets Hannah there. Nothing is at stake for Jill, and she knows nothing of Hannah's situation. Wondering whether Hannah will be there, Jill utters to a friend, 'Well, Hannah was at the bank two weeks ago on a Saturday. So she knows the bank will be open on Saturday.'

Ignorant High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. But neither Hannah nor Sarah is aware of the impending bill, nor of the paucity of available funds. Looking at the lines, Hannah says to Sarah, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning.'

High Attributor–Low Subject Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah calls up Bill on her cell phone, and asks Bill whether the bank will be open on Saturday. Bill replies by telling Hannah, 'Well, I was there two weeks ago on a Saturday, and it was open.' After reporting the discussion to Sarah, Hannah concludes that, since banks do occasionally change their hours, 'Bill doesn't really know that the bank will be open on Saturday'.

Suppose that, in all five situations, the bank will be open on Saturday. Here, I take it, are the intuitive reactions we have about these cases. In *Low Stakes*, our reaction is that Hannah is right; her utterance of 'I know the bank will be open' is true. In *High Stakes*, our reaction is that Hannah is also right. Her utterance of 'I don't know that the bank will be open' is true. In *Low Attributor–High Subject Stakes*, our intuition is that Jill's utterance of 'she knows the bank will be open on Saturday' is false. In *Ignorant High Stakes*, our reaction is that Hannah's utterance of 'I know the bank will be open tomorrow' is false. In *High Attributor–Low Subject Stakes*, our reaction is that Hannah's

utterance of ‘Bill doesn’t really know that the bank will be open on Saturday’ is true.

The *practical facts* about a situation are facts about the costs of being right or wrong about one’s beliefs. All five cases involve people with the same non-practical basis for the belief the bank will be open the next morning (in the first four, Hannah, and in the fifth, Bill). But the facts as to whether the relevant attributor can truly ascribe the predicate ‘knows that the bank will be open’ to the relevant subject vary. Furthermore, the facts vary in accord with the *importance* to some person—either the knowledge attributor or the putative knower—of the bank’s being open. This provides a prima facie case for the thesis that knowledge is not just a matter of non-practical facts, but is also a matter of *how much is at stake*.

I will call the thesis that knowledge does not depend upon practical facts *intellectualism*.³ Intellectualism is a wide orthodoxy. So conservatism demands the exploration of alternative paths. For example, one might attempt to explain away the force of the intuitions behind these scenarios, by arguing that, when someone recognizes that the costs of being wrong are particularly high, his or her confidence is shaken. The result of having one’s confidence shaken is either to reduce one’s degree of belief below the threshold required for knowledge or to defeat the evidence one has for one’s belief in some other manner. This explanation provides an elegant account of the second scenario, where Hannah’s awareness of the costs of being wrong undermines her confidence in her belief.⁴

However, this line of defense falters when one considers Ignorant High Stakes. In this case, Hannah’s confidence that the bank will be open is not shaken, because she is ignorant of the potential costs of not depositing her check. So the defender

³ Thanks to Earl Conee for suggesting this term.

⁴ Jon Kvanvig (on the blog *Certain Doubts*) suggested this as an account of these sorts of cases.

of this line of defense would have to adopt the position that Hannah does not know that the bank will be open in the second scenario, but *does* know that the bank will be open in the fourth scenario. And this is an odd position. After all, Hannah is more knowledgeable about her situation in the second scenario than she is in the fourth scenario. It does not seem correct that adding a little ignorance increases knowledge. In short, if Hannah does not know in the second scenario, it seems she also does not know in the fourth scenario. If so, then appealing to loss of confidence does not help in evading the consequence that practical interests can have epistemic consequences.

This line of defense also does not account for our intuitions concerning High Attributor–Low Subject Stakes. We may suppose that Bill’s confidence that the bank will be open is not affected by Hannah and Sarah’s situation. So the account does not provide an explanation of our intuition that Hannah and Sarah are correct to deny knowledge to Bill. So some other explanation is required.

Another strategy that proponents of intellectualism commonly appeal to in the face of these examples is to argue that in certain cases our responses are sensitive not to whether the subject knows, but to whether the subject *knows* that she knows. According to advocates of this strategy, Hannah knows that the bank will be open in Low Stakes, High Stakes, and Ignorant High Stakes, and Bill knows that the bank will be open in High Attributor–Low Subject Stakes. Our judgments to the contrary in the latter three cases are to be explained by the fact that the relevant subjects do not *know that they know* in any of these cases. According to this line of reasoning, knowing that one knows that *p* requires having more evidence for *p* than knowing that *p*. When we are aware that the stakes are particularly high for a subject, we tend to require not just that the subject knows the propositions upon which she bases her actions, but that she knows that she knows those propositions. Our awareness of the raised stakes for Hannah in High Stakes leads us to think that she needs to

know that she knows that the bank will be open, and not merely know that the bank will be open. Since she does not face a potentially hazardous predicament in Low Stakes, we are not led to make the error of thinking that she does not know that the bank will be open.

I am inclined to reject the KK thesis that knowing that p entails knowing that one knows that p . But I have difficulty seeing how the falsity of that thesis can be brought to bear to explain away these intuitions. First, the proponent of this way of rejecting our intuitions about these cases must explain why the fact that an agent does not know that she knows that p would lead us to deny that the agent knows that p . This requires an entirely independent explanation. Secondly, the proponent of this response must give some good reason to believe that in each case in which someone in a 'low-stakes' situation (such as Hannah in Low Stakes) seems to know that p , whereas someone with comparable evidence in a 'high-stakes' situation does not seem to know that p , the person in the low-stakes situation does not know that she knows that p .

I am skeptical that a good justification for the second claim can be provided. Most ordinary assertions of knowledge are made on such a basis that we can envisage someone in a higher-stakes situation (often a much higher-stakes situation), whom we would not think of as possessing that knowledge, given similar evidence. The proponent of this response would have to argue that, in all such cases, the person in the low-stakes situation knows that p , but does not know that she knows that p . This leads to widespread failure of knowledge of knowledge. It is one thing to deny that knowledge entails knowing that one knows, but it is quite another to license such a wholesale denial of knowledge of knowledge.⁵

⁵ Here is another point against the knowledge of knowledge maneuver, due to unpublished work by Kripke. Suppose that Hannah, in the low-stakes bank case, knows that the bank will be open. Suppose Bill has the same evidence as

A third reaction one might have when confronted by these cases is to explain them away as various types of *framing effects*, of the sort familiar from recent psychological studies of rationality. It has been established that our judgments about the rationality of various inferences are highly dependent upon idiosyncratic facts about how the background situation is described. It would be unwise to put very much weight upon this evidence in claims about the nature of rationality. Similarly, one might think that the intuitions we have in the above cases are also due to psychological framing effects. If so, they are unlikely to be helpful in inquiry into the nature of the knowledge relation.

However, the above cases reveal intuitions that are not analogous to the framing effects we see in ordinary speakers' judgments about rationality. The latter sort of judgment does not follow a discernible pattern that reflects any plausible general claim about rationality. In contrast, the intuitions we have in the above cases are just the intuitions we would expect to have, if certain antecedently plausible conceptual connections between knowledge and practical reasoning were true. As other anti-intellectuals have argued (Fantl and McGrath 2002, and especially Hawthorne 2004), it is immensely plausible to take knowledge to be constitutively connected to action, in the sense that *one should act only on what one knows*.⁶ For various theoretical reasons, this

Hannah, and is also in a low-stakes situation. Then Bill can felicitously and truly utter the sentence 'I know that Hannah knows that the bank will be open'. It seems bizarre to hold, as the advocate of this maneuver must, that Bill knows that Hannah knows that the bank will be open, but Hannah does not know that Hannah knows that the bank will be open, despite the fact that they have the same evidence that the bank will be open.

⁶ John Hawthorne (2004: 30) puts the principle as 'one ought only to use that which one knows as a premise in one's deliberations', which is a good way to elucidate the relevant sense of 'act on'. Hawthorne writes, concerning this principle: "There are complications that call for *ceteris paribus* style qualifications. 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

immensely plausible claim has not traditionally been accepted by those studying practical reasoning. But rejecting this claim devalues the role of knowledge in our ordinary conceptual scheme.

A standard use of knowledge attributions is to justify action. When I am asked why I went to the store on the left, rather than the store on the right, I will respond by saying that I knew that the store on the left had the newspaper I wanted, but I did not know whether the store on the right did. When my wife asks me why I turned left rather than going straight, I reply that I knew that it was the shortest direction to the restaurant. When it turns out that it was not a way to go to the restaurant at all, my wife will point out that I only *believed* that it was the shortest way to the restaurant. To say that an action is only based on a belief is to criticize that action for not living up to an expected norm; to say that an action is based on knowledge is to declare that the action has met the expected norm.

The fact that knowledge is thus connected to action is obscured by several points. First, *assertion* is also conceptually connected to knowledge; asserting that *p* implicates that one knows that *p*. So, in defending an action based upon one's knowledge that *p*, it is enough simply to assert that *p*. Secondly, in certain special circumstances, we do occasionally act on our knowledge that there is a *chance* that *p*, rather than our knowledge that *p*.⁷ For example, there are lotteries in which it is rational for me to buy a ticket, even though I do not know that I will win; when pressed to defend my purchase, I will respond that there is a chance I will win. But this is just to say that there are certain types of actions that I perform on the basis of beliefs about chances. In order for these actions to be acceptable, such beliefs must still constitute knowledge.

quickly.⁷ But *ceteris paribus* style qualifications are needed only insofar as they are needed in all normative claims. A similar point holds for the knowledge rule for assertion, discussed below.

⁷ Thanks to Jim Pryor for discussion here.

The intuitions we have in the above cases are best explained by appeal to our commitment to the principle that one should act only upon what one knows. For example, in High Stakes, we think it is mistaken for Hannah to act on her belief that the bank will be open on Saturday, and wait until Saturday to go there. The obvious reason why Hannah should not wait until Saturday to go to the bank is that she does not know that the bank will be open. The same is true for Ignorant High Stakes. Indeed, the intuitions in virtually all of the above cases are exactly the ones we would expect to have if it is true that knowledge is connected to action in the above sense.⁸ The intuitions therefore provide powerful intuitive evidence for an antecedently plausible principle concerning the relation between knowledge and action.

It is odd to assert instances of the schema ‘*P*, but I don’t know that *P*’ (Moore’s Paradox). The oddity of asserting instances of Moore’s Paradox is often taken to be strong evidence for the intuitive connection between assertion and knowledge (e.g. Williamson 2000: 253–5), that *one ought only to assert what one knows*. It is highly unlikely that the oddity of Moore’s Paradox is due to a psychological framing effect. For a similar reason, the reactions we have to virtually all of the cases I have discussed are not random noise. They are rather natural reflections of the conceptual connections between knowledge and action, of our intuitive adherence to the principle that one should act only upon what one knows.

So there is no easy intellectualist strategy for explaining away the intuitions. This leaves the intellectual with the following quandary. If the thesis that one’s knowledge of one of one’s true beliefs depends only upon non-practical facts is correct, then it cannot both be the case that (for example) Hannah

⁸ I say ‘virtually all the cases’, because the one intuition that remains mysterious from this perspective is the intuition we have in High Attributor–Low Subject Stakes. It is fine for the person in Low Stakes to act on his or her belief that the bank will be open.

knows that the bank will be open in Low Stakes, and does not know that the bank is open in the other three relevant situations. For, by stipulation, the non-practical facts for Hannah are the same in all of these cases, and she even has the same degree of confidence in her belief (at least in Low Stakes and Ignorant High Stakes). So, either the thesis must be rejected, or some other natural assumption.

Here are the options available to one who wishes to preserve the independence of knowledge from practical facts:

- (a) One can challenge the claim that these are the intuitions we have in these cases.
- (b) One can reject the semantic significance of one of the intuitions. For example, one could deny semantic significance to the intuition that the proposition semantically expressed by Hannah's utterance in Low Stakes is true. Alternatively, one could deny semantic significance to the intuition that the proposition semantically expressed by Hannah's utterance in High Stakes is true (or reject the semantic significance of either of the intuitions in the other two cases).
- (c) One can deny that the proposition expressed by Hannah's utterance in Low Stakes is really the denial of the proposition expressed by Hannah's utterance in High Stakes (and make similar maneuvers for the other two cases).

Though I certainly do not take all of the intuitions we have in the above cases as indefeasible, I will not discuss except in passing the first of these options. The role of these intuitions is not akin to the role of observational data for a scientific theory. The intuitions are instead intended to reveal the powerful intuitive sway of the thesis that knowledge is the basis for action. Someone who denies that we have many of these intuitions is denying the pull of the link between knowledge to action. But the *value* of knowledge is explicable in part by its links to action; it is for this reason that skepticism threatens agency. Those who deny these

intuitions are in effect maintaining that some other notion, such as appropriately confident belief, is intuitively the genuinely valuable one. It is because I find this reaction so implausible that I will not seriously consider rejecting these intuitions. Nevertheless, while my central interest is to evaluate accounts that make as much sense of these intuitions as possible, the central claims of this book hold, even if some of the above intuitions are less robust than others. I will leave it to the reader to decide which arguments in the book are strengthened or weakened by her particular pattern of intuitions.

As far as the second of these options is concerned, the most obvious way to develop it is to appeal to a certain view about the relation between semantics and pragmatics. According to this view, our intuitions about what is said by utterances of sentences are not in general reliable guides to the semantic contents of sentences in context, even relative to perfectly clear hypothetical circumstances like the ones described above. On this view, our intuitions about what is said by a sentence are often influenced by pragmatic, post-semantic content conveyed by the act of asserting that sentence.

For example, one might argue that we are wrong to think that Hannah's utterance in *Low Stakes* expresses a true proposition, because 'know' expresses a relation that holds between a person and only a very few select propositions, those for which (say) she has deductive valid arguments from a priori premises. But knowledge ascriptions may pragmatically convey that the subject stands in some epistemically looser relation with the proposition. One could then 'explain' the mistaken intuition on the hypothesis that we often confuse what an assertion of a sentence pragmatically conveys with the semantic content of that sentence relative to a context.

Giving pragmatic explanations of apparently semantic intuitions is a standard maneuver in philosophy. While this strategy is certainly occasionally called for, it must be applied with great circumspection. For example, DeRose (1999) considers a crazed

theorist who defends the view that 'bachelor' just expresses the property of being a man. This theorist holds that the intuition that 'is a bachelor' cannot be truly predicated of a married man has no semantic significance; it is due rather to (say) pragmatic felicity conditions governing the use of the term 'bachelor'. DeRose's point in considering such examples is that the tendency philosophers have to give pragmatic rather than semantic explanations of apparently semantic intuitions threatens to undermine the whole enterprise of giving semantic explanations. As he writes (1999: 198), concerning pragmatic explanations of speakers' apparently semantic intuitions about the cases that motivate his favored view:

It's an instance of a general scheme that, if allowed, could be used to far too easily explain away the counterexamples marshaled against any theory about the truth-conditions of sentence forms in natural language. Whenever you face an apparent counterexample—where your theory says that what seems false is true, or when it says that what seems true is false—you can very easily just ascribe the apparent truth (falsehood) to the warranted (unwarranted) assertability of the sentence in the circumstances problematic to your theory. If we allow such maneuvers, we'll completely lose our ability to profitably test theories against examples.

By undermining the data for semantic theory, this kind of strategy threatens to undermine the semantic project.

Of course, there are cases in which it is legitimate to provide pragmatic explanations of apparent semantic intuitions. Again, to borrow an example from DeRose (1999: 196 ff.), if someone clearly knows that p , it seems extremely odd to say that p is epistemically possible for that person. But there is a clear explanation from Gricean principles for the oddity in question. There is a general conversational principle to the effect that one should always assert the most informative proposition one is in a position to assert. If x asserts 'It is possible that p ', then x implicates, via this maxim, 'I do not know that p '. Our sense that such an

assertion is odd, or seems false, is due to the fact that x is implicating something known to be false. The problem with many pragmatic explanations of apparently semantic intuitions is that there is no such clear explanation from general conversational principles.

Denying the semantic significance of apparently semantic intuitions is a significant cost, one that we should be reluctant to bear in the absence of a clear explanation of these intuitions from general conversational principles. Since I am not aware of such an explanation, I think that the most fruitful way to pursue preserving intellectualism is by appeal to the third option. And this leads us to the thesis of contextualism.