

Precis of Knowledge and Practical Interests

Jason Stanley

Our intuitions about whether someone knows that *p* vary even fixing the intuitively epistemic features of that person's situation. Sometimes they vary with features of our own situation, and sometimes they vary with features of the putative knower's situation. If the putative knower is in a risky situation and her belief that *p* is pivotal in achieving a positive outcome of one of the actions available to her, or avoiding a negative one, we often feel she must be in a particularly good epistemic position to know that *p*. If however she is not in a risky situation, we tend to be considerably more permissive in our attributions of knowledge. Our intuitions about when someone knows that *p* and when she doesn't also fluctuate with what is *salient*, either to the putative knower or ourselves. If someone claims to know where their car is, and I raise the possibility that it has been stolen, she may retract her claim to know. Which of these tendencies, if any, reveal something about the nature of our epistemic terms or concepts?

Call the *practical facts* of a person's situation those facts that bear on the costs and benefits of the actions at that person's disposal. Does knowing depend upon practical facts? Some of the intuitions we have about the conditions under which someone knows suggest that knowing does depend upon practical facts. Given the role the concept of knowledge has in our conceptual scheme, this should not be surprising. It is standard to use the verb "know" in the appraisal of action; it is an old wife's adage that it is dangerous to act without knowing all the facts. The value of knowledge lies in part in the fact that it is the mental state that ought to guide us in our actions. If knowledge and action are conceptually linked, then it would be no surprise that knowing something is related to being rationally permitted to act upon a belief. Since the latter is a function of practical facts, it would be no surprise to discover that the former is as well. The connection between knowledge and action therefore suggests that whether someone knows that *p* at a time *t* depends at least in part upon practical facts. If so, then knowing is what I call an *interest-relative* relation. The thesis of the book is that knowing is an interest-relative relation.

Raising something to salience has the effect of making new interpretations available for context-sensitive expressions. Thus, if we take seriously the effects of

salience on our judgments about when agents know and when they do not, we will likely end up as *contextualists* about knowledge claims. Contextualism has a number of putative benefits above and beyond a charitable explanation of many of our intuitions. For instance, it promises to deliver a response to skepticism that both accounts for the force skeptical arguments have on us, as well as our disinclination to believe that they undermine ordinary claims to know. It can provide such an explanation, consistent with preservation of single-premise epistemic closure. Contextualism also can explain many of the cases in which we think someone does not know something because of heightened risk, by assimilating such cases to ones in which counter-possibilities to the proposition putatively known are more salient to the knowledge attributor. Contextualism can provide such an explanation without undermining the attractive thesis of *intellectualism*, which is that whether or not someone bears an epistemic property to a proposition does not depend on the practical facts of their situation. Finally, contextualism in epistemology fits with a several decade old trend in philosophy generally of finding contextualist solutions to long-standing philosophical problems, such as the sorites paradox, the liar paradox, and the problem of material constitution. Contextualism and interest-relativity are not in and of themselves competing accounts; one can be a contextualist, and hold that the various different denoted knowledge relations are interest-relative in nature. But if one is a contextualist, one can provide explanations for many of the phenomena that lead one to reject intellectualism. Since intellectualism is a powerful and perhaps non-negotiable theoretical commitment of many philosophers, the desire to preserve it is a central motivation for contextualism.

That salience affects our intuitions about knowledge attributions should not lead us *directly* to contextualism. Raising certain features of a situation to salience also may significantly distort our judgments. The tendency we have to give up knowledge-attributions in the face of salient counter-possibilities may be a psychological framing effect rather than a semantically significant feature of our epistemic practice. Since salience is a dubious nail on which to hang a theory, other sort of corroboration should be available in support of the contextualist thesis in epistemology. Since contextualism is an all-purpose strategy for evading philosophical problems (albeit the favored strategy of some of the best and the brightest philosophers), there is extra reason for those not yet

convinced of the contextualist thesis in epistemology to seek more evidence in support of its central claims.

In Knowledge and Practical Interests (henceforth K&PI), I spend three chapters evaluating the costs and benefits of contextualism in epistemology. There are several reasons for spending so much space rebutting the contextualist thesis. First, contextualism is the most plausible alternative to the interest-relativity of knowledge for many of the intuitions we have about the effects of practical facts upon the truth of various knowledge ascriptions. Secondly, by focusing intensively on the one case of contextualism in epistemology, I show what sorts of considerations can be advanced for and against a contextualist thesis about a kind of discourse. This kind of investigation is particularly important, since, as anyone familiar with non-historical areas of philosophy in the past three decades knows, contextualism is a widely employed strategy in addressing philosophical problems. Finally, since my own work in philosophy of language is characterized by sympathy for non-obvious context-sensitivity in natural language, I am naturally predisposed towards the doctrine.

In the fifth chapter of K&PI, I develop a detailed version of an interest-relative theory of knowledge. Following Williamson (2000), I reject reductive analyses of the knowledge relation. The form epistemology should take is to state facts about the primitive knowledge relation, and that is my strategy in this book. My purpose in developing an analysis of knowledge along interest-relative lines was not to show that an analysis of knowledge is possible, but merely to show that interest-relativity does not add unmanageable complexity to the epistemological project. Furthermore, by engaging in the fiction of analyzing the knowledge-relation along interest-relative lines, I am able to investigate some of the central notions involved in investigating the interest-relativity of knowledge, such as the notion of *serious practical question* (a proposition that a rational agent must take seriously in decision making).

In the sixth chapter, I turn to an explicit comparison between the thesis of interest-relativity and the thesis of contextualism as explanations for various puzzling epistemological problems, including the problem of skepticism. It is tempting to think that the reason we do not need to worry about skepticism is that skeptical propositions are not serious practical questions for us. But I do not think that that the fact that knowing is

an interest-relative relation provides a viable solution to the problem of skepticism. This does not favor contextualism, since I also do not think that contextualism provides a viable solution to the problem of skepticism. At the very least, to the extent it does, its solution can be mimicked by the advocate of the interest-relativity of knowledge.

Interest-relativity is not a special methodology for evading philosophical problems. Rather, it is a way of using the philosophical problems to gain insight about the metaphysics of the properties and relations in question. In contrast, contextualism is a special methodology that is applicable everywhere. Another special methodology is *relativism about truth*. Like contextualism, relativism is a method that can be applied to any philosophical dispute in any area. In contrast to contextualism, which is obviously correct about many apparent disputes, it is not clear that there is a single apparent dispute for which relativism about truth is the resolution. Nevertheless, relativism about truth has been fashionable in recent years, thanks in large part to its impressive development and defense in the work of John MacFarlane. In the seventh chapter, I turn to a critique of relativism about truth as applied to apparent disputes about knowledge. I spend less time on relativism, for several reasons. First, unlike contextualism, it has very little empirical support from uncontroversial stretches of discourse. Secondly, if one is concerned to defend the “purity” of the knowledge relation from invidious infection by practical facts, adopting relativism about truth is scarcely the method to do so. Those unsympathetic with the view that knowledge is a matter of practical facts are unlikely to be assuaged by the prospect that the truth-value of propositions containing the knowledge relation depends upon such facts, though its metaphysical analysis does not. Those who are uncomfortable with accepting the intrusion of practical facts into the analysis of the knowledge relation are likely to be even more uncomfortable with their intrusion into the analysis of truth.

My main purpose in K&PI is to elaborate and defend the thesis of interest-relativity about knowing. But my subsidiary purpose is methodological. Resolutions to philosophical problems have increasingly taken the form of applications of general methodologies, applicable everywhere (e.g. contextualism, fictionalism, relativism). This is a worrisome development, especially when there are few agreed upon methods of evaluating the application of such global methodologies. In writing K&PI, I wanted to provide rigorous introductions to some of these methodologies, and discussions of the

sorts of considerations that are relevant for and against them. In chapter 8, I therefore broaden my outlook to consider these issues against the philosophical backdrop of quandaries in areas outside epistemology. My focus in Chapter 8 is on contextualist and quasi-interest-relative theories about vagueness, though my considerations bear on contextualist views about the liar paradox as well. I conclude that there is little plausibility in appealing to an interest-relative metaphysics of properties in these domains. The fact that knowing is an interest-relative relation is an important epistemological discovery, one that helps elucidate the connections between knowledge and action. But interest-relativity does not seem to help extricate us from any of the classic philosophical quandaries.

Replies to Gilbert Harman, Ram Neta, and Stephen Schiffer

Reply to Harman

In “Epistemic Contextualism as a Theory of Primary Speaker Meaning”, Gilbert Harman endorses a version of contextualism that takes the primary speaker’s meaning, or utterer’s meaning, of a sentence containing an instance of “know that p” to vary with epistemically relevant features of the context, such as what the speaker takes for granted in the conversation in which the utterance takes place. But, as I argue below, this construal of contextualism does not help Harman evade the detailed criticisms of contextualist positions I give in K&PI. Nevertheless, I conclude by pointing out that, given his general epistemological commitments, Harman is in fact in a better position to defend epistemic contextualism from some of my detailed criticisms than its other proponents.

Why does Harman believe that his version of contextualism, in terms of primary speaker’s meaning, evades the detailed criticisms of contextualism given in K&PI? Many of the criticisms I give in K&PI have to do with the mechanism that mediates between the sentence uttered and the proposition it expresses. For example, I argue at length that no uncontroversial context-sensitive expressions behave in the way that contextualists say that instances of “know that p” behave. There is therefore no evidence of the existence of the sort of mechanisms required by the contextualists about knowledge-ascriptions

evident in our use of language outside such constructions. This should lead us to be skeptical of the existence of such mechanisms in the case of knowledge ascriptions.

Harman seems to think that there is a view of semantic theorizing that is silent about the mechanism by which an utterance of a sentence gives rise to its primary speaker meaning. As he writes:

Importantly, epistemic contextualism understood in this way is not directly a claim about the meanings of words or sentences, although one or another theorist might go on to try to explain the supposed differences in primary speaker meaning by appeal to something special about the meaning of the word “know” or other aspects of the meaning of sentences used to make knowledge attributions.

But this quote does not provide us with some new and distinctive view of semantic theorizing. Rather, it expresses Harman’s *neutrality* about the mechanism that mediates between knowledge-ascriptions and their context-sensitive primary speaker meanings. Harman seems to think that neutrality about the mechanism that mediates between sentence uttered and proposition expressed helps him to evade my criticisms. But my criticisms were intended to show that there is *no* plausible mechanism that supports contextualism about knowledge-ascriptions. Any mechanism, whether semantic or pragmatic in nature, that does the work that the contextualist in epistemology requires, must have certain features (for example, it must not allow mid-discourse context-shifts). If the relationship between sentences containing “know” and their primary speaker meanings were governed by such a mechanism, then we would expect to see its workings elsewhere in our use of language. But we do not. It is difficult to see why failure to commit to any particular account of the mechanism that mediates between sentences and the propositions their users employ them to express allows Harman to evade the criticism that there is no such mechanism. To respond to the criticisms, one would need to provide a plausible mechanism, rather than insist on contextualism anyway, and maintain neutrality about how the contextualist apparatus functions.

Here is a perhaps extreme analogy. Suppose that someone wanted to argue that the primary speaker meaning expressed by “Snow is white” was the proposition that grass is green. One might object to such a theorist that there is no way, given compositionality of meaning, and the meanings of “snow”, “is”, and “white”, that the

primary speaker meaning of “snow is white” could be the proposition that grass is green. If they wish to maintain that the primary speaker meaning of “snow is white” is the proposition that grass is green, they must commit to “snow” meaning something other than snow, or “is” meaning something other than predication, or “white” meaning something other than the property of whiteness. It hardly constitutes a sufficient defense of the view, in the face of this kind of criticism, to deny that one is “making a claim about words or sentences” in saying that the primary speaker meaning of “snow is white” is the proposition that grass is green. That is, neutrality about semantic mechanisms is not an option. Rather, defending the view would require one to give some plausible explanation of how a sentence that contains words that appear to have a semantic content that determines one proposition could be regularly used by speakers to express something that does not seem related to that semantic content at all.

Harman claims that speakers use one and the same knowledge ascription to express different propositions relative to different contexts of use. He does not reply to my criticism that there is no plausible way to implement this proposal. I have already stated my reasons why it is not enough to make a claim about speaker meaning, in the face of the kind of criticisms I provided. But many of the criticisms of contextualism I give in K&PI are independent of the mechanism by which the context-sensitivity is brought about, and are focused specifically on the claim that one and the same knowledge ascription can be used to express different propositions relative to different contexts of use, by whatever mechanism. For example, the contextualist can explain our intuitions in a case in which the person attributing knowledge is in a high-stakes situation (i.e. a lot depends upon whether the attributor is right), and the person to whom knowledge is being attributed is in a low-stakes situation (nothing much depends upon whether she is right); call this a “High Attributor-Low-Subject Stakes” situation. But the contextualist cannot explain our intuitions about an example that looks completely parallel to this one (K&PI, p. 115):

High Attributor-Low Subject Stakes*

It is Friday, and Bill is considering going to the bank. Nothing much hangs on whether he gets to the bank before Monday. So he considers delaying going to the bank until Saturday. He has been to the bank before on Saturday, and found it open. So he assumes it will be open the next day (and it will). He utters absent-

mindedly, ‘Well, since I know the bank will be open tomorrow, I’ll just go tomorrow.’ In contrast, it is very important for Hannah to get to the bank before Monday. She overhears Bill, and asks him what the basis for his claim was. But she too has been to the bank on previous Saturdays, which does not allay her concern that the bank has changed its hours. She therefore concludes that what Bill said was false.

Bizarrely, the contextualist concludes that while Hannah makes no mistakes in High-Attributor-Low Subject Stakes, she does make an error here, since what Bill said was in fact true. Yet there is no structural or intuitive difference between these cases. The objection is not that there are cases for which the contextualist cannot provide a non-error theoretic explanation; we are all in that situation. It is rather that there are pairs of cases that are exactly structurally parallel, and about which we have exactly similar reactions, but the contextualist is forced to give a charitable account of one, and an error-theoretic account of the other. Only contextualist theories have this kind of consequence, and Harman’s version of contextualism does not rescue him from this predicament.

So Harman has not provided, in his remarks, a sufficiently compelling defense of contextualism. However, I think that Harman is in a much better position than most other contextualists to defend contextualism. One of the primary motivations for contextualism, dating back to Gail Stine’s pioneering paper, was to reconcile certain kinds of responses to skepticism with single-premise epistemic closure. The contextualist’s thought is that a sentence such as (1) cannot truly be uttered:

- (1) John knows that he has hands, but he doesn’t know that he isn’t a bodiless brain in a vat

The reason, according to the contextualist, that (1) cannot truly be uttered is because single-premise epistemic closure is correct, and the word “know” cannot switch its content in mid-discourse. In *K&PI*, I argued that this motivation for contextualism is inconsistent with the semantic thesis of contextualism, since semantically context-sensitive expressions do generally behave in a way that would allow true and indeed felicitous utterances of this sentence, since context-sensitive expressions can change their content in mid-discourse. But Harman’s contextualism, as he makes clear in his contribution (and as anyone familiar with his epistemology knows), is not motivated by

adherence to single-premise epistemic closure. Indeed, as Harman notes, he thinks that sentences such as (1) express true propositions. So Harman's contextualism, unlike the contextualism of Stine, Lewis, and DeRose, is not one that is motivated chiefly by the desire to reject the truth of utterances of sentences such as (1). Mid-sentence context-shifts are perfectly compatible with Harman's desired version of contextualism. As a consequence, he has a response to some of the criticisms I level against contextualism in Chapter 4 of K&PI, one that is not available to its other defenders.

Reply to Neta

The initial paragraphs of Neta's contribution consist of advertisements for other work of his in which he purports to have refuted conclusions of mine. Since for the most part he does not provide his arguments, I cannot respond to most of his contentions (e.g. "...there is a version of intellectualist contextualism that provides a better explanation of the other four cases..."). But, since he does provide a glimpse into one of his argument strategies, I can hazard a response to at least that one here. In chapter 3 of K&PI, I argue that uncontroversial context-sensitive expressions allow for mid-sentence context-shifts. For example, in a sentence such as "Syracuse has many unemployed men and many serial murderers", the two occurrences of the context-sensitive word "many" express different determiner meanings (since what counts as "many" for a serial murderer population is distinct from what counts as "many" for employment purposes). But some of the central motivations for contextualism are inconsistent with the possibility that distinct instances of "know that p" can have different contents within a discourse. So the contextualist cannot simultaneously maintain that instances of "know that p" are context-sensitive, and that their context-sensitivity explains (say) the oddity of assertions of the falsity of single-premise epistemic closure.

Indexical expressions such as "I" and "here" allow mid-sentence context-shifts with difficulty. But this is not to any distinctive semantic fact, but rather due to the fact that it is odd for the speaker of an utterance to change during a discourse. Neta writes, "As I argue elsewhere...analogous mundane facts determine the same standard of knowledge throughout a discourse, and so Stanley's linguistic arguments against

indexicalist contextualism fail.” Without having any clue about what kinds of arguments would support the contention that it’s just as odd to change the standard of knowledge within a discourse as it is to change speakers, I will nevertheless grant this for the sake of argument. This would not show that my “linguistic arguments against indexicalist contextualism fail.” First, words like “I”, “here” and “now” are narrow indexicals, which are a small subspecies of context-sensitive expressions generally. I provide numerous arguments for the conclusion that the word “know” is not a narrow indexical. So even if “mundane facts” determine that a standard of knowledge must remain the same throughout the discourse that does not entail that my arguments against indexicalist contextualism fail. Finally, if “know” is not a narrow indexical, but rather is associated with a variable for epistemic standards, then even if mundane facts determine that epistemic standards cannot shift within a discourse, different occurrences of the same instance of “know that p” within a discourse could have different contents, because the variable for epistemic standards could be controlled by different operators. For example, in “Some man loves his mother, and a neighbor loves his mother too”, the two occurrences of “his” may have different interpretations, since their interpretations are controlled by the different quantifiers “every man” and “a neighbor”; something similar occurs with comparison class variables associated with adjectives such as “tall” and “large” (K&PI, p. 58).¹ Similarly, if there are variables for epistemic standards, then distinct variables may be bound by different operators, consistently with the contextually salient epistemic standard remaining fixed.

Let me now turn to the arguments Neta explicitly provides, rather than the ones to which he merely alludes. Neta’s first point is that IRI entails that someone with the same evidence for p and q might know that p, yet not know that q, because of her greater investment in the truth of p. Neta is right to emphasize that this kind of consequence is to be avoided, if possible. However, it is only a problem for certain versions of IRI. The purpose of my book was to defend IRI about knowledge, and I was officially neutral on IRI about other epistemic notions, such as evidence. But as I say in the conclusion of

¹ The operator controlling the interpretation of the pronominal elements (“his” as well as the comparison class variables associated with adjectives) is actually most likely to be a lambda-abtractor associated with the verb phrase.

K&PI (p. 182), "...my own view is that all epistemic notions are interest-relative." Neta's discussion provides further support for the view that interest-relativity affects all epistemic notions.

Suppose it is a serious practical question for me whether I am on Main Street, so serious that I cannot rely on my friend Bill's assurance that I am on the corner of State and Main. In that case, Bill's testimony does not provide the kind of evidence it would if I had less of a practical stake in some proposition I actually base upon it. I do not formulate the thesis of the interest-relativity of evidence in K&PI. But a way of formulating it is suggested by Neta's discussion of methods. The quality of evidence for a person X at a time t provided by some information, or some method of gaining information, depends upon whether it is used to support X's belief in a proposition that is a serious practical question for X at t. So, Bill's testimony is not good enough evidence for me to know that I am on either State or Main, since it is not very good evidence, given my practical situation.

The central objection in Neta's paper is not that I should accept the thesis of the interest-relativity of evidence. It is rather that the resulting view is in serious tension with the Knowledge-Action Principle (KAP), the thesis that *one should act only on what one knows*. Neta's main point is that even if Kate doesn't know that she will soon die if she does not find Main Street, she is rational to search for Main Street frantically. As Neta recognizes, the advocate of KAP has a simple reply; by appeal to knowledge of probabilities. Neta's response is:

Can she nonetheless reasonably search for Main Street on the premise that there is a *chance* that she will soon die if she does not find Main Street? Perhaps, but, given her background knowledge, she could just as easily have known *that* premise about chances even before the jogger's assertion – knowledge of human mortality would suffice.

But this is not the proposition knowledge of which would legitimate her actions. Rather, Kate's search for Main Street would be legitimate, provided she knew there was a *sufficient* chance that she will die soon if she does not find Main Street, where what constitutes a sufficient chance is determined (say) by rational expected utility. So the advocate of KAP and IRI has a simple response to Neta's concern.

However, I am not committed to the constitutive relation between knowledge and action taking the precise form of KAP. I endorsed KAP in K&PI, because I thought it was the most likely candidate to be the principle that reflects the intuitive normative connection between knowledge and action. I now think it was a mistake to emphasize KAP to the exclusion of other possible connections between knowledge and action that would be sufficient to derive IRI. For example, it is possible to obtain IRI from the considerably less contentious premise that knowing that *p* is *sufficient* for acting on one's belief that *p*, together with some intuitions about cases. For example, in Low Stakes, Hannah knows that the bank will be open, and can act on this knowledge. Intuitively, in High Stakes Hannah cannot act on her belief that the bank will be open. Therefore, by the claim that knowing that *p* is sufficient for acting on *p*, Hannah does not know that the bank will be open in High Stakes. Thus, IRI follows from some intuitions about cases, together with the far less controversial claim that if one knows that *p*, it is rationally permissible to act on one's belief that *p*. I do not think that the claim that knowledge is sufficient for action exhausts the connections between knowledge and practical reasoning; more needs to be added to it to explain the role knowledge plays in the appraisal of action (see Hawthorne and Stanley, forthcoming). But if all one cares about is deriving IRI, appeal to the sufficiency claim is sufficient. None of the considerations raised by Neta are relevant to this way of deriving IRI. But then again, Neta has not succeeded in undermining the more contentious argument for IRI based upon KAP.

Reply to Schiffer

According to the standard way of thinking of the relation between knowledge and action, knowledge is a flawed folk notion, and we may be systematically wrong about its application. An account of what makes a decision rational should not appeal to it, not the least because whether or not a decision an agent makes is rational should be independent of the truth or falsity of skepticism. In Schiffer's challenging contribution, "Interest-Relative Invariantism", he develops a series of worries for the motivation for Interest-Relative Invariantism I provided in K&PI, ones which are informed by this classical way of looking at the relation between knowledge and action.

Before I turn to a systematic discussion of Schiffer's objections, I must explain why I did not provide a reductive analysis of knowledge along interest-relative lines.

Schiffer writes:

...IRI leaves plenty of room for IRI theorists to disagree among themselves about what IRI account of knowledge is correct. Stanley doesn't commit to any particular IRI account. He does consider one account, but he does so merely "for the sake of explicitness" (88) and "just for the purposes of illustration" (89), and he acknowledges that the account is neither complete nor correct.

There is a perhaps unintended implicature here to the effect that the fact that I was unable to provide a complete, correct, reductive account of knowledge along interest-relative lines already suggests dubiousness about the thesis of the interest-relativity of knowledge. But the reason I did not provide a complete and correct account of knowledge is that, following Williamson (2000), I don't believe that there is such an account of knowledge along *any* lines, including interest-relative ones. In light of this fact, the way epistemology should be conducted is by trying to state various facts about the (irreducible) knowledge relation, and that is the strategy employed in K&PI.

I now turn to responding to Schiffer's main concerns. In K&PI, I suggest that the norm that links knowing and rationally acting is that there is a necessary condition on acting on a belief that the belief is an instance of knowledge (the Knowledge Action Principle, or KAP). Schiffer provides a number of arguments against KAP. I am in fact sympathetic to some of the concerns Schiffer raises about KAP. My concern with Schiffer's discussion is that he fails to acknowledge any of the pressures that would lead one to seek some norm relating knowledge to action. By going so aggressively on the offense to undermine the letter of what I said, he blinds himself to the motivations that led me to that position. My strongest commitment is to the claim that knowledge is connected to practical rationality, and its value lies in part in this connection. I am less committed to this connection taking the form of KAP.

I appealed to KAP for the following two reasons. First, the intuitions in the stakes examples suggest that our views about when someone knows and when it is proper for them to act are intertwined. Secondly, as I emphasize in K&PI (p. 10), "A standard use of knowledge attributions is to justify action". These facts are utterly inexplicable by appeal

to the standard Bayesian account of rational action in terms of expected utility. Indeed, part of the very motivation of the Bayesian account is to eliminate reference to knowledge in an account of action. In contrast, if we do adhere to a norm of practical rationality that links rationally acting on a belief to knowing, then the facts to which I draw attention in K&PI are explicable. But there are a number of other ways of explaining the link between knowledge and rational action besides KAP (see Hawthorne and Stanley (forthcoming)). The options are therefore not just between the standard Bayesian account of rational action, and KAP. Be that as it may, I will defend KAP in what follows.

Schiffer's objections to KAP center around his adherence to the standard Bayesian account of rational action, and his conviction that it is often rational to act on partial beliefs. In K&PI, I accounted for such cases by arguing that cases in which it seems to be rational to act on a partial belief in p are really cases in which it is rational to act, because one is acting on one's belief that there is a chance that p , or that is it sufficiently likely that p . Schiffer complains that assertions such as "There is a chance of rain", or "It probably will snow" are expressions of the asserter's degree of belief, rather than reports of propositions about epistemic probabilities. I doubt that this is correct, since I think that when one is reporting on one's beliefs about the weather, one intends to assert propositions about objective rather than epistemic chances. In cases in which it is utterly clear that it is epistemic chance that is at issue, Schiffer's intuitions are somewhat less compelling. Suppose, watching a basketball game on television that I know is on tape-delay, and unaware of the outcome, I assert that the team losing by 10 points has a good chance of winning nevertheless, since their best three point shooter has yet to get his characteristic hot hand. But unknown to me, it turns out that their best three point shooter was injured. If someone tells me after this information is announced that I was wrong, I can easily reply by asserting "Well, I didn't know that, so for all I knew, they were going to win". If my original use of "good chance" was epistemic, what this amounts to is a reaffirmation of the truth of my original utterance, suggesting that my original statement was a report of a proposition about epistemic probability after all.

Schiffer writes that "Stanley doesn't say what is wrong with the Bayesian account". What is wrong with the Bayesian account is that it renders completely

mysterious the fact that we standardly use “know” to appraise actions. An account of rational action is unsatisfactory when it renders it completely mysterious why we use the concepts we do in appraising actions. Schiffer misses this point when he complains that:

Perhaps, however, the most telling objection to the F-M-H line is that requiring knowledge of epistemic probabilities that warrant partial beliefs is such a retreat from the original idea and motivation behind the idea that “one should act only on what one knows” that it would seem to amount to nothing more than a so far unmotivated qualification of the Bayesian expected-desirability account.

What motivates the qualification of the Bayesian expected-desirability account is the need to explain what that account manifestly does not, namely our tendency to employ the concept of knowledge in appraising the rationality of actions.

Schiffer objects that, in a standard case of apparent rational action on the basis of partial belief, there is no additional need for the requirement that one knows that one is justified to some degree. He envisages Jane, who has evidence that warrants her belief that it will rain to degree .4, and writes “Surely, given her desires, her carrying an umbrella is justified if she’s justified in believing to degree .4 that it will rain tonight, whether or not she *knows* that there’s a non-negligible chance that it will.” That depends upon our intuitions about a case in which Jane is justified in believing to degree .4 that it will rain tonight, but does not in fact know that there is a possibility it will rain. Schiffer tries to force our intuitions in his favor by appealing to our intuitions about someone who “lacks the concept of the relevant epistemic probability and doesn’t even believe, let alone know, that there’s a ‘non-negligible chance’ that it will rain.” But it is not easy to conceive of someone who is not capable of having the relevant belief. One can believe that there is a chance of rain, without having the “concept of the relevant epistemic probability”, because in attributing the belief that there is a chance of rain, the chance-predicate is *de re* with respect to the belief (in other words, to attribute belief that there is a chance that *p* is to attribute the belief, about a chance, that the person believes the proposition has *that*). No sophisticated conceptual apparatus is required to grasp propositions of that sort.

I am nonetheless sympathetic to some of Schiffer’s concerns. It may be that KAP is a norm of rational action, but it might be a *defeasible* norm, one that can be trumped by

other norms, such as expediency. There may be multiple norms governing rational action, which sometimes conflict. It is therefore consistent with KAP that one sometimes acts on beliefs that are not knowledge, because there are other rational pressures on our behavior. Finally, as I have mentioned, there are a number of other potential norms linking knowledge to action that are consistent with rational action on partial belief, yet also entail IRI (though Schiffer is of course not responsible for discussing them). For example, as I discussed in my reply to Neta, if knowing that p is merely *sufficient* for acting on p , one can derive the interest-relativity of knowledge. My commitment is to the thesis that our practice of employing the concept of knowledge to appraise actions is to be taken seriously in an account of both knowledge and action. However that thesis is to be reflected in an account of the norms of rational behavior, it will reveal the incompleteness of the standard Bayesian account of rational action, which makes no mention of knowing.

Schiffer's second criticism, due to Jeremy Landen, is that the intuitions that drive the stake examples are "a lot more fragile" than I take them to be. More specifically, one can engender other intuitions by appealing to cases in which someone raises a counter-possibility to a belief to salience in a forceful manner (as in the case attributed to Landen). But I have always been fully aware that one could engender other intuitions by bringing counter-possibilities to salience. The question is whether the intuitions engendered by the salience of counter-possibilities were reliable guides to the nature of knowledge or the semantics of "know".

Salience is typically associated with semantic context-sensitivity; the semantic content of a context-sensitive term is typically determined by facts about contextual saliency. If contextualism were true, it would therefore provide an elegant charitable account of how our intuitions about when a knowledge attribution is true co-vary with what is salient in the context of its use. But given that contextualism is false, it is difficult to see how the salience of counter-possibilities could be anything other than the familiar framing effects so much discussed in psychology. When I bring to salience the fact that Carol was a philosophy major who is single, bright, and was an anti-war activist, people find it more likely that Carol is a bank teller and a feminist than that Carol is a bank teller. Similarly, when I raise to salience a counter-possibility to p , that leads people to

retract their claims to know that *p*. Just as “[i]t would be unwise to put very much weight upon this evidence in claims about the nature of rationality” (K&PI, p. 9), so it would be unwise to place very much weight upon evidence about salience in claims about the nature of knowledge.

In contrast, our intuitions about the epistemic significance of stakes are not easily relegated to mere framing effects. The link we detect in our folk discussions of action between the concept of knowledge and practical rationality accords smoothly with the intuitions we have about the stakes examples. The stakes examples are other manifestations of the same tendency that leads us to criticize someone who acts on her belief that *p*, for acting without knowing that *p*. It is because I take this tendency to be a central rather than a peripheral feature of our concept of knowledge that I reject the possibility that our intuitions about the stakes examples are framing effects. The connection between knowledge and practical rationality is in large part what gives knowledge the value it has in our conceptual scheme. Since there is no value whatsoever in having a single relation that we bear to propositions that comes and goes with what possibilities are salient, I see no reason for a non-contextualist to treat the effects of salience on our judgments about knowledge as anything other than framing effects. As we know from the literature on framing effects, their relative robustness does not constitute such a reason.

Schiffer’s third section concerns a series of cases he calls the “smirk” examples. In these cases, someone asserts that they know a proposition, and their assertion is countered with a *recherché* counter-possibility. In response, they “smirk”. Schiffer’s error-theoretic account of these cases is that the original assertions of knowledge were false, but conversationally appropriate. Though any account of knowledge is committed to error theoretic explanations of some class of data, Schiffer’s discussion of his smirk examples suggests a broad-sweeping error theory that is unacceptable to anyone other than an epistemological skeptic. For Schiffer’s smirk examples are no different in kind from cases such as:

A walks into C’s living room, carrying a heavy plate of food. A’s vision of the living room is obscured by B, who is standing in front of her.

A: Do you know whether there is a table here?

B: I do know; I'm looking right at one.

A: But how do you know it's not merely a hologram of a table?

B: Smirks.

The above case is just like Schiffer's other smirk examples. If one shares Schiffer's view that the smirk cases are ones that demonstrate the original claim to knowledge was false but conversationally appropriate, then one would also conclude that B's claim to know is false but conversationally appropriate. In short, to take Schiffer's smirk cases at face-value is to accept external-world skepticism. In my book, discussions of external world skepticism are purely negative – I emphasize that this or that kind of view does not help answer the skeptic. The purpose of my book is not to reply to external world skepticism, but rather to explore the options available to the non-skeptical epistemologist.

Schiffer's final worry concerns a potential conflict between his smirk examples and IRI. Schiffer writes:

Some smirk examples will also be cases that one might think are problematic for IRI in a way Stanley recognizes, but whose prima facie problematic nature he attempts to explain away. Here is one such example. The proposition that Mavis is visiting her sister isn't a serious practical question for her husband Alvin. Let's suppose that its epistemic possibility on Alvin's total evidence is great enough for him to know that Mavis is visiting her sister by the lower standards that low stakes makes sufficient for knowledge. At the same time, the proposition that Mavis wasn't killed in a car accident on the way to her sister's place *is* a serious practical question for Alvin, and by the lights of IRI, his evidence for that high stakes proposition isn't good enough for him to know it. Nevertheless, it's apt to seem that *pace* IRI, he *is* in a position to know that high stakes proposition by the same low standards by which he knows that Mavis is visiting his sister. For, it would seem, he can simply reason "Mavis is visiting her sister, and that entails that she wasn't killed in an accident on the way to her sister's; so Mavis, thank god, wasn't killed in a car accident."

Schiffer then argues that to block this kind of derivation, I would appeal to the fact that x's drawing an inference q from a premise p can lead to a change in x's practical situation. Whereas p was not a serious practical question before the inference was drawn, once the serious practical question q is inferred from p, p becomes a serious practical question. As Schiffer writes, "I suppose the idea is that a proposition that is low stakes

will become high stakes when one infers a high stakes proposition from it.” Schiffer then writes:

But if that is the idea, it’s hard to see how Stanley can accept it. For, as his deployment of the example Ignorant High Stakes shows, he acknowledges that a proposition can be high stakes for x even though x isn’t aware that it is. Given this, it shouldn’t be x’s *inferring* a high stakes proposition from p which makes p high stakes; it should be x’s *entailing* a high stakes proposition, whether x knows it or not. But Stanley also can’t accept that, for if every proposition that entails a high stakes proposition is itself high stakes, then there will be virtually no low stakes propositions.

Schiffer is correct that I certainly do not think that a high stakes proposition is one that entails a high stakes proposition. But his discussion also betrays confusion about the central notions of my developed example.

I never speak of propositions being high stakes or low stakes for a person; this notion has no place in my theory. Rather, I speak of propositions being *serious practical questions* for a person at a time. This is a notion I borrow from a brief discussion by James Joyce of what he calls a *serious practical possibility*. As Joyce writes (2002, p. 90), “H is a serious practical possibility for DM if and only if she is rationally required to factor the possibility of H’s truth into her decision making.” I make various suggestions about how to cash out this notion, e.g. with the use of warranted expect utility.

Human beings are not ideal rational agents; we have limited capacities. Given this fact, there are certain questions we are rationally required to resolve before acting, and other questions we are not rationally required to resolve. We would view someone as neurotic who checked to see if there was a banana in her tailpipe whenever she went to start her car. But we would view her actions as perfectly reasonable if there were neighborhood boys who regularly stuck bananas in car tailpipes in her neighborhood, and she was aware of this fact. In the former situation, the proposition that there is a banana in the tailpipe of her car is not a serious practical question for her at the time of her checking, whereas in the latter situation, it is a serious practical question for her.

A serious practical question is a proposition the possible truth of which a person with limited resources is required to factor into her decision making in order to meet the demands of rationality in that particular situation. Though it is extremely difficult to come up with a precise account of this notion, beings with limited resources are certainly

not rationally required to factor into our decision making all of the logical consequences of our beliefs. The idea of a serious practical question has its home in the study of the epistemic responsibilities of agents with limited time and evidence. The existence of propositions that are serious practical questions for an agent, even though that agent is not aware that they are, is nevertheless entirely consistent with the foregoing. That is, I assume that any good explication of the notion of serious practical question will have as a consequence that there will be propositions the truth of which an agent is rationally required to factor into her decision making, though she is not aware that she should. This does not entail, or even vaguely suggest, that all entailments of her beliefs are ones the agent must factor into her decision making. That would be quite inconsistent with the notion of a serious practical question.

I do think that in many cases, drawing an inference may alter one's practical situation in such a manner as to make propositions that previously were not serious practical questions into serious practical questions. This may occur in a number of ways. But Schiffer's example is not one of them. The proposition that Mavis wasn't killed in a car accident is not a serious practical question in the scenario envisaged by Schiffer, unless (for example) Alvin has suddenly been offered life insurance for that day for Mavis at a particularly good rate (in which case the proposition that Mavis is visiting her sister is a serious practical question as well). The proposition that Mavis wasn't killed in a car accident is not a serious practical question for Alvin, because only a neurotic would factor that into their decision making; he is not rationally obligated to factor it into his decision making at that time.

Schiffer ends his contribution by mentioning that IRI has the counterintuitive consequence that one can know more by caring less (I suppose that Schiffer would prefer a surgeon who was closely related to him, rather than one had less of a personal stake in the matter). This objection would be more compelling if we were negotiating a space of alternatives which included ones that lacked any starkly counter-intuitive consequences.

Bibliography

Hawthorne, John and Stanley, Jason (forthcoming): “Knowledge and Action”

Joyce, James (2002): “Levi on Causal Decision Theory and the Possibility of Predicting one’s own Actions”, *Philosophical Studies* **110**: 69-102.

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