
THE JOURNAL OF PHILOSOPHY

VOLUME CV, NO. 10, OCTOBER 2008

KNOWLEDGE AND ACTION*

What is the relation between knowledge and action? According to one standard picture, there is none. Rational action is a matter of maximizing expected utility, where expected utility is a function of utility and subjective credence. It is subjective degrees of belief that matter for rational action, not knowledge. On this picture, having knowledge that p is independent of whether it is rational to act on one's belief that p : knowledge that p is not sufficient since one may know that p despite lacking sufficiently high subjective credence to warrant acting on the proposition that p ; and knowledge that p is not necessary, since high subjective credence can rationalize action even in the absence of knowledge.

In contrast to the picture just sketched, our ordinary folk appraisals of the behavior of others suggest that the concept of knowledge is intimately intertwined with the rationality of action. Suppose, for example, that Hannah and Sarah are trying to find a restaurant, at which they have time-limited reservations. Instead of asking someone for directions, Hannah goes on her hunch that the restaurant is down a street on the left. After walking for some amount of time, it becomes quite clear that they went down the wrong street. A natural way for Sarah to point out that Hannah made the wrong decision is to say, "You shouldn't have gone down this street, since you did not know that the restaurant was here." Here is a similar case. Suppose John decides not to buy health insurance anymore, reasoning that he is healthy enough.

* The authors of this paper are listed in alphabetical order. We have given this paper at the Rutgers/Princeton Graduate Conference, Oxford University, University of Edinburgh, St. Andrews University, University of Calgary, Harvard University, and the University of Aberdeen. We are grateful to all the members of the audience who participated at those occasions. Keith DeRose, James Pryor, and Stephen Schiffer gave us helpful comments on written versions of the paper. Finally, as usual, the greatest thanks are due to Timothy Williamson, who provided both oral and written commentary at various stages throughout the process.

He calls his mother to report excitedly on his money-saving decision. His mother can berate him for not buying the insurance, by appealing to the fact that he does not know that he will not fall ill. Another case is as follows. You are offered a cent for a lottery ticket that cost a dollar, in a 10,000 ticket lottery with a \$5,000 first prize and reason as follows:

I will lose the lottery

If I keep the ticket I will get nothing

If I sell the ticket, I will get a cent

So I ought to sell the ticket.

This piece of practical reasoning is absurd. It is not acceptable to act on one's belief that one will lose the lottery (or one's belief that if one keeps the ticket, one will get nothing). The most natural explanation for why one cannot act on these beliefs is that these beliefs are not knowledge.

Consider also how knowledge interacts with conditional orders. Suppose a prison guard is ordered to shoot a prisoner if and only if they are trying to escape. If the guard knows someone is trying to escape and yet does not shoot he will be held accountable. Suppose meanwhile he does not know that someone is trying to escape but shoots them anyway, acting on a belief grounded in a baseless hunch that they were trying to escape. Here again the person will be faulted, even if the person is in fact trying to escape. Our common practice is to require knowledge of the antecedent of a conditional order in order to discharge it. The lesson transfers to instructions with less weighty consequences than death by gunshot. Suppose a sous-chef is told by the master chef to take a cake out of the oven if it is done before the master chef returns. It will cast doubt on his competence were he to take the cake out without knowing whether it is done—whether or not the cake is in fact done. There is something of an analogy here with the knowledge-assertion connection. The kind of defect manifested by the sous-chef who takes acts on that order without knowing whether the cake is done is similar to the kind of defect manifested by a sous-chef who answers 'yes' to the question 'Is the cake done?' when he does not know whether it is.

Consider next how blame, judgments of negligence, and so on, interact with knowledge. If a parent allows a child to play near a dog and does not know whether the dog would bite the child, and if a doctor uses a needle that he did not know to be safe, then they are *prima facie* negligent. Neither the parent nor doctor will get off the hook by pointing out that the dog did not in fact bite the child and the needle turned out to be safe, nor by pointing out that they were very

confident that the dog/needle was safe. Of course, some excuses are acceptable but these too are sensitive to the facts about knowledge. If the parent knew that they did not know that the dog would bite the child, and if the doctor knew that he did not know that the needle was safe, we will deem the action inexcusable. If such second order knowledge is absent we will be more open to excuses. Here again there is an analogy with the case of assertion. If someone asserts that p without knowing it and knowing that they do not know that p , they will have no excuse for their failure to adhere to the norm that one should assert only if one knows. If on the other hand, they assert that p , do not know that p , but cannot be expected to know that they do not know that p , we may be willing to deem their failure to comply with the norm excusable. The conceptual structure, one familiar from the normative realm, explains suitable appraisal in terms of a combination of norms and excuses for failure to comply with them.

Consider finally certain occasions when beliefs are properly disregarded. Suppose someone is sufficiently paranoid to believe that his hands are dirty even after he has washed them, knows this about himself, and forces himself to disregard that belief in his conduct. It seems crucial to the propriety of his behavior that he knows that the belief is not a piece of knowledge, which in turn leads to a recognition that the content of that belief is not a suitable basis for action.

Many epistemologists would argue that justified belief is enough for the propriety of the relevant behavior in many of the examples we have given. We will have much to say in response to this view in the pages that follow. For now, it bears emphasis that (in English at least) it is considerably more natural to appraise behavior with the verb 'know' than the phrase 'justified belief', or even 'reasonable belief'. Perhaps this is because 'know' is a phrase of colloquial English, whereas 'justified belief' is a phrase from philosophy classrooms. But this is itself a fact that should be surprising, if the fundamental concept of appraisal were justification rather than knowledge. Finally, there are cases in which no appeal to justified belief will do, not even as an excuse. For example, suppose Hannah's husband Mordechai has gone off to war, and goes missing for many years. Hannah remarries after waiting five years, reasonably assuming her husband to be dead. After reemerging from captivity, Mordechai might legitimately complain to Hannah that she should not have remarried without knowing that he had died. It is reasonable for Mordechai not to be satisfied with the excuse that Hannah had a justified belief that he was dead.¹

¹ Thanks to Jim Pryor for the example.

Judging by our folk appraisals, then, knowledge and action are intimately related. The theories of rational action with which we are familiar leave this unexplained. Moreover, discussions of knowledge are frequently silent about this connection. This is a shame, since if there is such a connection it would seem to constitute one of the most fundamental roles for knowledge. Our purpose in this paper is to rectify this lacuna, by exploring ways in which knowing something is related to rationally acting upon it, defending one particular proposal against anticipated objections.

I. A PREVIOUS PROPOSAL

In their carefully constructed paper "Evidence, Pragmatics, and Justification," Jeremy Fantl and Matthew McGrath² assume the following principle about the relation between knowledge and action:

The Knowledge-Action Principle. If you know that p , then it should not be a problem to act as if p .

They recognize that their use of 'as if' requires clarification. After all, it may be rationally mandatory to pretend to others that not- P even when one knows that P ; under such circumstances, there is a perfectly good sense in which it is not all right to act as if P . Their idea is that if one knows p during a period then the optimal ranking of states of affairs matches the ones that are optimal conditional on P .³ Thus:

The Refined Knowledge Action Principle

If X knows P during period d , then for any choice between states of affairs x_1, \dots, x_n during d , X is rational to prefer one state of affairs A to another state of affairs B iff X is rational to prefer A to B conditional on P .⁴

They then apply this principle to particular cases in accordance with their favored decision theoretic framework. That framework is one

² Fantl and McGrath, "Evidence, Pragmatics, and Justification," *The Philosophical Review*, cxi, 1 (2002): 67–94.

³ Where P has a nonzero credence, A is preferable to B conditional on the proposition that P iff A and P is preferable to B and P . Where P is zero then one's conditional preferences are undefined or else we need to introduce primitive conditional preferences that cannot be defined in terms of outright preferences. This theoretical choice is similar to that concerning conditional belief. Where Q is zero we either say that beliefs conditional on Q are all undefined or else introduce primitive conditional beliefs. Those, like us, who opt for a framework in which knowledge brings with it probability 1 ought to take seriously the introduction of primitive conditional preferences and beliefs.

⁴ The relevant notion of preferability, standard in decision theory, is one according to which A is preferable to B only if one welcomes the news that A to the news that B ("welcome the news" is a somewhat technical notion here, because in practice receiving the news that A typically involves receiving further information as well that might affect the preference ranking).

according to which a probability assignment to P that is less than one does not automatically shift upwards when knowledge is acquired, and, in consequence, one according to which knowledge that P is compatible with a probability assignment of less than one. Suppose P is .7 for S at t_1 , but for some choice at t_1 , a state of affairs that is preferable on P is not preferable *simpliciter*. By the principle above, P is not known. Suppose no new evidence is acquired during the interval t_1 to t_2 but at t_2 the test articulated by the Knowledge Action Principle is passed. As the authors are conceiving of things, so long as P is true and other knowledge-friendly conditions are in place, the subject knows P at t_2 but not at t_1 , with no change in evidential probability.⁵

We do not wish to contest the authors' central principle. But we wish to voice three concerns. First, we do not think that it provides sufficient illumination on how knowledge connects with action. Second, the surrounding theoretical framework seems to get things backwards as far as the order of explanation is concerned. Third, that framework makes trouble for a plausible closure principle concerning knowledge. Let us speak to each of these points in turn.

There are many cases where acting on P does not lead to untoward actions but where we intuitively judge that one ought not to treat P as a reason for action. Suppose for example that one is offered five dollars for a lottery ticket whose expected return is only two dollars. Suppose one sells it on the grounds that one will lose. Here the action is correct, but it is not acceptable to use the proposition that one will lose as a reason for acting. The judgment remains even if it is in fact true that one will lose. Similarly, suppose there is a significant risk that an operation will kill someone but that it is a risk worth taking and suppose further that a doctor performs the operation, reasoning that the patient will not die. The action is preferable, but the reason for acting is unacceptable. Yet the principle does not explain this. The state of affairs in which the surgery is performed is preferable in fact to the state of affairs in which the surgery is not performed, and the state of affairs in which the surgery is performed is preferable to the state of affairs in which the surgery is not performed, conditional on the proposition that the patient will survive the surgery. Intuitively speaking, the Knowledge-Action Principle captures the idea that it is acceptable to use that which one knows as a basis for action. But it is silent on whether it is acceptable to use that which one does

⁵ In fact, Fantl and McGrath formulate a principle in terms of *justification* rather than knowledge. But, as they admit, the motivation for their principle governing justification comes ultimately from a corresponding principle about knowledge.

not know as a basis for action. But this misses out on a good deal of the important structural ties between knowledge and action.

Turning to the second point, consider a case where the relative preferability of a pair of choices conditional on the proposition that p does not match their relative preferability *simpliciter*. The Knowledge-Action principle tells us—correctly in our view—that one does not know that p . Yet the decision theoretic gloss that Fantl and McGrath offer gets the order of explanation the wrong way around. Consider the case where one is offered a penny for a lottery ticket, or where on Friday one considers postponing a visit to the bank until Saturday in a “high stakes” situation, where much depends upon reaching the bank before Monday. The envisaged account tells us that in situations such as these, one does not know the relevant proposition and that this is because it would be irrational to act on that proposition. But that pattern of explanation does not fit our intuitive diagnosis. Our reaction to the lottery situation is not that the subject does not know that he will not win the lottery because it is not rational to act upon that supposition. It is rather that it is not rational to act on the proposition that she will not win because she does not know that she will not win. Similarly, our reaction to a high stakes bank case is not that someone does not know on Friday that the bank will be open on Saturday, because it is not rational for that person to wait until Saturday. It is rather that it is not rational for the person to wait until Saturday to go to the bank, because she does not know that the bank will be open.

A further concern about the theoretical framework in which Fantl and McGrath are operating is that it makes trouble for multi-premise closure, which is the principle that one can generate knowledge by deduction from sets of premises that one knows. Suppose I am faced with a choice between A and B . B does not risk a fine. If law x is in force then if I do A , I pay a 10 pound fine to authority w , and if law y is in force then if I do A I pay a 10 pound fine to authority v . I do not prefer to do A conditional on having to pay both fines. Suppose I am .9 confident that law x is in force and .9 confident that law y is in force but only .8 confident that both are in force. I prefer doing A to doing B , since the risk of two fines is sufficiently low. It is not at all hard to find utility assignments such that it is true that I prefer that I do A to that I do B iff I prefer that I do A to that I do B conditional on law x being in force, and I prefer that I do A to that I do B iff I prefer that I do A to that I do B conditional on law y being in force, but I do not prefer that I do A to that I do B iff I prefer that I do A to that I do B conditional on law x being in force and law y being in force. Suppose further that credence .8 is sufficient for full belief, and one has deduced the conjunction that x is in force and law y is in force

from the conjuncts. Fantl and McGrath's framework militates in favor of the conclusion that in this setting (supposing nothing else is pertinent), one knows that law x is in force, one knows that law y is in force, but does not know that law x and law y are in force, and this despite firmly believing that $p \& q$ on the basis of the belief in the conjuncts. We are aware that many contemporary epistemologists are more than ready to jettison multi-premise closure as an unrealistic fantasy. In our view this may have far-reaching consequences that have not been properly appreciated. For now, at least, let us underscore the conditional claim that if one holds multi-premise closure in high esteem, one ought to be wary of Fantl and McGrath's theoretical framework.

In short, despite its intuitive appeal, the Revised Knowledge-Action Principle and the surrounding theoretical baggage that Fantl and McGrath provide do not provide an adequate account of the connections between our knowledge and our judgments about practical rationality.

II. OUR PROPOSAL

In many cases in which we use 'know' in the appraisal of action, someone acts upon some practical reasoning she employs, and is thereby subject to criticism. The criticism takes the form of the charge that the agent in question did not know one of the propositions upon which she acted. The norm of practical rationality most *immediately* suggested by this behavior, and one which we shall defend, is the following:

The Action-Knowledge Principle

Treat the proposition that p as a reason for acting only if you know that p .

The Action-Knowledge Principle makes immediate sense of our use of 'know' to criticize the actions of others. When someone acts on a belief that does not amount to knowledge, she violates the norm, and hence is subject to criticism. That is why we use epistemic vocabulary in criticizing the actions of others. (Of course, as noted earlier, we will in some cases be open to excuses where the norm is violated. It would a grave mistake to respond to this by complicating the norm: as in other realms, we should recognize that some violations of norms are more excusable than others.)

Note that an analogous principle seems plausible for reasons for belief, namely: Treat the proposition that p as a reason for believing q only if one knows that p . One attractive feature of the Action-Knowledge Principle is that it unifies the practical and theoretical domain of reasons: if it is correct, then proper reasons for belief and reasons for action have a uniform nature.

The Action-Knowledge Principle straightforwardly accounts for the use of epistemic terms in appraisals of practical reasoning. I cannot sell my lottery ticket for a penny on the basis that I will lose, since I do not know that I will lose the lottery. I am subject to criticism for driving to Syracuse without purchasing ultra cheap insurance in a situation where I reason 'I will not get into an accident', because I did not know that I would not get into an accident. The Action-Knowledge Principle is the most direct way of accommodating our intuitions about the relation between knowledge and action in these cases.

Knowing that p is necessary for treating the proposition that p as a reason for acting. Is knowing that p also sufficient for treating the proposition that p as a reason? Not quite. There are lots of cases where the question whether p is absolutely irrelevant to the case at hand. In that case it would be odd to say that it is appropriate to treat the proposition that p as a reason for acting, even if one knows that p . To get a plausible sufficient condition, we need some restriction to choices for which the proposition that p is relevant. Let us say that a choice between options $x_1...x_n$ is p dependent iff the most preferable of $x_1...x_n$ conditional on the proposition that p is not the same as the most preferable of $x_1...x_n$ conditional on the proposition that not- p . For p -dependent choices, it seems highly intuitive that knowledge that p makes it appropriate to treat p as a reason for action. We can now combine this idea with the Action-Knowledge Principle articulated earlier to get:

The Reason-Knowledge Principle

Where one's choice is p -dependent, it is appropriate to treat the proposition that p as a reason for acting iff you know that p .⁶

Our principle concerns what is *appropriate* to treat as a reason for action, rather than what one *ought* to treat as a reason for action. It would be overly demanding to require someone to treat *all* of their relevant knowledge as reasons for each action undertaken. The principle is therefore a claim about what is permissible to treat as reasons for action in a given choice situation.

⁶ Note that this makes the fact that there is an external world an appropriate reason for acting. Suppose I have a choice between trying to have sex and doing mathematics. Suppose conditional on there being an external world, I prefer the former, but conditional on there being no external world I prefer the latter. What this means is that quite often some very general facts that are appropriate to consider as reasons are rarely explicitly noticed as such.

We do not go so far as to say that when one's choice is not p -dependent it is never appropriate to treat the proposition that p as a reason for acting. For example, there are cases where the proposition that p may be a reason for doing X but is overridden by other reasons to such an extent that one's preference for not doing X is not contingent upon the proposition that p . Thanks to Selim Berker for discussion here.

Note that there is a good case to be made that in the primary sense of personal reasons—where a personal reason is picked out by a construction of the form ‘His reason for acting’—it is propositions that are personal reasons, and where the following principle holds: the proposition that p is one’s reason for acting only if it is known.⁷ And supposing such a principle is apt, one is tempted to reformulate the Action-Knowledge Principle as: Treat the proposition that p as a reason for acting only if it is a reason for acting.

Recalling the concerns that we voiced in connection with Fantl and McGrath, let us inquire further as to what kind of decision theoretical framework is the most suitable setting in which to embed this principle. It is clear enough that if we want multi-premise closure, we had better operate with a notion of probability according to which knowledge delivers probability 1.⁸ Moreover, within that setting, we get the order of explanation right in those cases we looked at earlier where we criticize action on the basis of an absence of knowledge. Suppose we criticize someone for selling a lottery ticket for next to nothing on the grounds that he did not know. Supposing we think of knowledge as delivering probability 1, we get a smooth explanation: to sell a lottery ticket for next to nothing, the probability that it would lose would have to be 1 or fantastically close to 1. But the only ready means for securing an epistemic space that fits that probabilistic criterion would

⁷ The qualification ‘primary sense’ is important. After all, there is a good deal of predication shiftiness possible in ‘His reason’. Consider: (1) Jason was his reason for leaving the party. (2) His reason for leaving the party was that it was boring. (3) His reason for leaving the party was that he thought that it was boring. (1) involves a use of ‘His reason’ that identifies reasons with objects. (2) involves a use for which it is plausible to apply the Reason-Knowledge Principle. (3) involves a use for which that principle seems less applicable (the truth of (3) does not obviously require that the subject knew that he thought that it was boring).

The claim that something is a reason for acting only if it is known has been defended by John Hyman (“How Knowledge Works,” *The Philosophical Quarterly*, XLIX, 197 (1999): 433–51) and Peter Unger (*Ignorance: A Case for Skepticism* (New York: Oxford, 1975)). As Hyman writes, “...the fact that p cannot be A’s reason for doing something if A does not know that p ” (p. 444). And in Unger (chapter 5), he argues at length for the following claim (his “third and final premiss” for his argument that skepticism leads to irrationality):

If S’s reason (for something X) is that p , then S knows that p ; and if S’s reasons (for X) are that p and that q and so on, then S knows that p and S knows that q and so on (*ibid.*, p. 200).

Hyman and Unger are defending the claim that something is a reason for you only if it is known. They do not defend the additional normative principle that something is appropriately treated as a reason only if it is a reason in their sense. The latter is what we are defending.

⁸ More generally, a piece of knowledge is at the top end of the scale of probabilities. We prescind from a full discussion of whether that scale is suitably represented by real numbers, which representation theorems are excessively idealized, and so on. For now, we assume that a real number scale is in order.

be by knowing that the ticket was going to lose. In effect, one points out that the difference between knowing that the ticket was going to lose and one's actual epistemic probability that that ticket was going to lose makes a crucial difference to the rationality of the action under consideration.

Note further that in a decision theoretic setting where knowledge delivers probability 1, Fantl and McGrath's Knowledge-Action principle falls out quite straightforwardly. Where the probability of p is not zero, A is preferable to B conditional on p iff A and p is preferable to B and p . Suppose one knows that p . Then in the current setting the proposition that p has probability 1. But then, quite obviously, A will be preferable to B iff A and p is preferable to B and p .

Thus ends our sketch of the tie between knowledge and action. By way of defense and further elaboration we turn to consider a series of objections and replies.

III. OBJECTIONS AND REPLIES

Objection 1. Given the decision theory described, why bother with all this talk of reasons? Why not just say that one ought to do that which delivers maximal expected utility and exploit epistemic probabilities (where, inter alia, knowledge deliver probability 1) as the ground of expected utility?

Reply. We are by no means opposed to a perspective according to which claims of practical rationality—and in particular what one ought to do—are grounded in a decision theory of the sort we have gestured at. But the need to integrate such a theory with reasons for action is still vital. For one thing, there are cases where one does what one ought to do but for the wrong reasons, and this phenomenon needs explanation. More generally we need to distinguish between the existence of a reason for acting and appreciating that reason in such as a way as to make it *your* reason for action (between mere rationalizers and motivators).⁹ As we are thinking about things, it is knowledge that constitutes the relevant sort of appreciation that converts the mere existence of a reason into a personal reason. Suppose for example, that your evidential probability that P is .5 and, as a result, one ought to prefer a certain contract $C1$ over another contract $C2$. In this situation, it is natural to say that the fact that one's evidential probability is .5 is a reason for accepting $C1$ over $C2$. But supposing evidential probabilities are not luminous, it is perfectly possible

⁹ Also, recalling earlier Davidsonian discussions, there are cases where one's action is rationally overdetermined, and we want to make sense of the question which of a range of adequate reasons for doing x were in actual fact motivating.

that in this situation one does not know that one's evidential probability that P is .5. In this situation, the fact that one's probability that P is .5 cannot function as one's reason for acting.¹⁰

Objection 2. It seems that we can rationally act on *partial beliefs*. Suppose I am driving to a restaurant, when I come upon a fork in the road. I think it is somewhat more likely that the restaurant is to the left than to the right. Given that these are my only options, and (say) I do not have the opportunity to make a phone call or check a map, it is practically rational for me to take the left fork. Yet I do not *know* that the restaurant is on the left. It seems that I am rational to act on partial beliefs that are not known.¹¹

Reply. Again, let us distinguish between what is rational to do, which is determined by one's favored version of expected utility theory, from what is appropriate to treat as one's reason for action. If (say) one's evidential probability that the restaurant is on the left is sufficiently high (say .7), then what it is rational to do may very well be to go left rather than go right. But, on our view, it is not proper to treat the proposition that the restaurant is on the left as a reason for going left. The envisaged objection to our view is that it is after all proper to act on the partial belief that the restaurant is on the left, even when one does not know that it is.

How would the advocate of the Action-Knowledge Principle and the Reason-Knowledge Principle account for cases in which it intuitively is rational to act on partial beliefs? One possible maneuver is to appeal to *knowledge of chances*:

...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.¹²

On this proposal, the normative role of partial beliefs is replaced by *beliefs about chances*. Many philosophers use 'chance' to refer to some

¹⁰ This gap between 'a reason' and 'your reason' is widespread and unsurprising: We allow that the fact that someone is unkind is a reason not to go out with them, even in a situation where this is unknown and so cannot count as your reason.

Of course it is likely that only a fragmentary part of the total decision theory story functions as a motivating reason in the sense described above. Very often one can know that if P , then x will be all things considered rational even when one is not in a position to know the full story about one's decision theoretic situation. There is plenty more to say here, but we shall not pursue the matter.

¹¹ This objection is forcefully made in Stephen Schiffer's "Interest Relative Invariantism," *Philosophical and Phenomenological Research*, LXXV, 1 (2007): 188–95.

¹² Jason Stanley, *Knowledge and Practical Interests* (New York: Oxford, 2005), p. 10.

kind of objective probability. In contrast to this, chance-talk in ordinary language is much more often epistemic in character. In accord with this, we will use 'chance' to express *epistemic probabilities*, where one's epistemic probability for a proposition is determined by *the total body of one's knowledge*.¹³

The intuitive description of instances of the problem of partial belief is in accord with our suggestion. For example, above we described a case in which someone believes it is more likely that the restaurant is to the left than to the right. This is obviously a belief about epistemic probability. One can treat this proposition about comparative likelihood as a reason for action if and only if one knows that it is more likely (or that it is sufficiently likely) that the restaurant is to the left than to the right. As with beliefs that are not about chances (that is, epistemic probabilities), there will of course be cases where one excusably acts on beliefs about chances that fall short of knowledge. As always, our diagnosis here will be that one violates the fundamental norm of practical reasoning but in an excusable way (not that every belief about epistemic probabilities is excusable). By contrast, cases where one acts on suitable knowledge of epistemic probabilities will be ones where the norm is not violated at all.

There are clear differences between cases in which we act upon our belief that there is a high probability that p , and cases in which we act on the belief that p . Suppose that I believe that there is a high probability that it will rain. On these grounds, I decide to bring an umbrella on my walk. This is a case in which I am treating my belief that there is a sufficient chance of rain as a reason for my action of taking my umbrella. But if I had believed that it *will* rain, I would have cancelled my walk.¹⁴ The proposal is certainly not that every time I act on the

¹³ Many readers might be more used to epistemic probability construed as the probability a proposition has for a person at a time, given that person's *evidence*. If, as Timothy Williamson has argued (in *Knowledge and Its Limits* (New York: Oxford, 2000), pp. 184–208), one's evidence just is one's knowledge, then epistemic probability in our sense and epistemic probability in this perhaps more familiar sense coincide. We are neutral about Williamson's $E=K$ thesis, but believe that many ordinary locutions involving probability express probabilities relative to one's total body of knowledge.

¹⁴ Thus, Ludwig Wittgenstein is wrong when he writes in paragraph 338 of *On Certainty* (G.E.M. Anscombe and G.H. von Wright, eds. (New York: Harper and Row, 1969)): "But imagine people who were never quite certain of these things, but said that they were *very* probably so, and that it did not pay to doubt them. Such a person, then, would say in my situation: 'It is extremely unlikely that I have ever been on the moon', and so on, and so on. *How* would the life of these people differ from ours? For there *are* people who say that it is merely extremely probable that water over a fire will boil and not freeze, and that therefore strictly speaking what we consider impossible is only improbable. What difference does this make in their lives? Isn't it just that they talk rather more about certain things than the rest of us?"

belief that p , I am really acting on the belief that there is a high probability that p . If I believe it will rain, I will act differently than if I believe that it is very likely that it will rain. In ordinary action, we treat many propositions that are not about chances as reasons for acting. The point is just that it is sometimes appropriate to treat propositions about chances as reasons for action.

Needless to say, in those cases in which we appropriately act on our knowledge of epistemic probabilities, it is not merely propositions to the effect that there is some chance that P that figure as reasons. Our knowledge of the level of likelihood is typically crucial. My reason for taking an umbrella on my walk is not that I know that there is a chance of rain. It is rather my knowledge that it is *reasonably likely* that it will rain, that is, that it is likely enough to warrant the disutility of bringing an umbrella on my walk.

One might worry that the current line conflicts with certain of the data with which we began.¹⁵ If knowledge of epistemic probabilities is sometimes an appropriate basis for action, why was it appropriate to criticize Hannah for doing what she did in the absence of knowledge of the location of the restaurant? In response, let us emphasize two points. First, it is clear enough that in certain cases, knowledge that p is reasonably likely is insufficient to warrant an action that would be warranted by knowledge that p . It is, after all, easy to imagine utility assignments on which it will be decision-theoretically rational to stop to ask directions in a situation in which one merely knows there is a pretty good chance that the restaurant is in a certain location. Second, there are many circumstances in which it is clear that the person will not have high evidential probability that p unless they know that p . That is, in a large range of ordinary cases, we reasonably think that if the evidential probability of some proposition for some person is not 1, it must be considerably less than 1. There is, after all, in ordinary life no easy technique for getting oneself in a situation where one does not know where the restaurant is but nevertheless acquires evidence that makes the evidential probability for the restaurant's being in a certain location very close to 1. Putting these points together, there will be plenty of situations where we judge that unless someone knows where the restaurant is, the rational thing to do will be to stop to ask directions—and that judgment will not be at all shaken by the concession that the protagonist knows that it is reasonably likely that the restaurant is in a certain place. Moreover, it is very natural to construe the scenario with which we began as a scenario of that sort.

¹⁵ Thanks to Susanna Siegel for pressing this objection.

Another worry with the thesis that one sometimes acts on knowledge of chances is that it seems to forge an implausible distinction between having p as a reason for action (in a case where the epistemic chance of p is 1) and having a fact about the chance of p as a reason for action (in a case where the chance of p is less than one). In the former case, one acts on knowledge that is just about the world, for example, the knowledge that the restaurant is on left. In the latter case, one seemingly acts on self-knowledge, for example, the knowledge that one's knowledge makes it reasonably likely that the restaurant is on the left. One might worry that this results in a schizophrenic account of the nature of personal reasons.¹⁶

This objection might have bite, if our proposal was that we sometimes act on our knowledge of *subjective* chances. But our proposal is rather that we sometimes appropriately act on our knowledge of *epistemic* chances, where the epistemic probability of a proposition for a person is determined by that person's knowledge. So, for example, that the proposition that the restaurant is on the left has a .6 epistemic chance for me might be determined, for example, by the fact that I know that 3 out of the 5 people I have asked for directions have indicated that it is in that direction. When I act upon my knowledge that it is reasonably likely that the restaurant is on the left, it would also be apt for me to act upon my knowledge of those propositions that *make* for a high epistemic chance of the proposition that the restaurant is on the left. Thus for example, when confronted with a fork in the road, I might turn left because (a) it is epistemically more likely that the restaurant is on the left *or instead* because (b) three out of five people told me that the restaurant is on the left. Acting on the knowledge that three out of five people told me that the restaurant is on the left is not acting upon knowledge of chances.

Our proposal concerns *what is appropriate to use as a reason for acting*. Whenever we appropriately act on our knowledge of the high epistemic chance of the proposition that p , we could equally appropriately have acted on knowledge of propositions that are not about chances, namely, those propositions we know that make for a high epistemic chance of the proposition that p . Our proposal vindicates ordinary usage in accommodating the fact that propositions about chances can be appropriate reasons for action. But whenever someone appropriately uses such a proposition as a reason for acting, there are propositions they know that are *not* about chances that they could instead have appropriately used as reasons for acting. In short, given the tight

¹⁶ Thanks to David Chalmers for pressing this objection.

connection between worldly knowledge and epistemic probability, the account that we recommend is far from schizophrenic.

Objection 3. Animals, children and perhaps some adults who do not have the concept of chance are nevertheless capable of correctly acting for reasons. Presumably, they are also capable of rationally acting on partial beliefs. But it does not seem on your account that they should have this capacity. For, by hypothesis, they are not capable of grasping propositions about chances.¹⁷

Reply. Given the previous point, this objection is irrelevant. As we emphasized above, when someone appropriately acts on a belief about epistemic chances, there are always propositions that are not about chances that they could instead appropriately use as reasons for action. Our proposal is therefore fully consistent with the possibility of beings who could act for reasons while lacking the concept of chance. (This is not to say that we agree that there are such creatures. As Daniel Dennett and others have noticed, one can predict and explain the patterns of behavior of all sorts of mundane artifacts by ascribing partial beliefs and desires to it. We would not want to take this as showing that we are entitled to describe such artifacts as possessing reasons for action, and for which the question of what it ought to treat as a reason is a proper topic.)

Be that as it may, one should not overintellectualize the concept of chance. One does not need to grasp a philosophical analysis of chance in order to have the concept of chance. Chance-like constructions ('it's likely', 'it's probable', 'there is enough of a chance', 'it was sufficiently likely') are pervasive in natural language. This suggests that there are very basic concepts of chance that we employ, though we are of course mostly ignorant of its nature (in this sense, of course, 'chance' is no different from 'water'). Someone does not need to have taken a confirmation theory class in order to grasp what is said by an utterance of 'There is a decent chance of rain so we should bring an umbrella'.

Objection 4. Suppose one is less than subjectively certain that *P* but knows *P*. Should not one in that case treat *P* as having less than probability 1 in the sense of probability relevant to rational action?

Reply. We hear this objection quite often but do not find it particularly compelling. Suppose someone knew that they had turned the coffee pot off, but having left the house are a little bit anxious about whether it is off. It is far from clear that we should craft our normative theory of action to deliver the conclusion that one ought in such a circumstance to go back and check. In general, it should be noted that intuitions go a little hazy in any situation that some candidate

¹⁷ We owe this objection again to Schiffer (*op. cit.*, p. 190).

normative theory says is sufficient to make it that one ought to *F* but where, in the case described, one does not know that situation obtains. As Timothy Williamson has emphasized, cases of this sort will arise whatever one's normative theory, given that no conditions are luminous. It arises even for the subjective Bayesian who says subjective credences are the guide to action. Suppose, for example, that one is subjectively certain that one has turned the coffee point on, but one has non-negligible credence that one is not subjectively certain. Here we do not have straightforward intuitions about what one should do. Similarly, adapting a discussion in Williamson, suppose one goes slowly from feeling cold to not feeling cold and some point is the last point where one is subjectively certain that one feels cold. The subjective Bayesian says that one should bet on feeling cold at any odds at that point but not after. But insofar as one is unable to know of oneself that that point is a point where one is subjectively certain, one will be a little uncomfortable and hesitant about accepting that claim that, even though one does not know one is subjectively certain, one should bet on feeling cold at any odds so long as it in fact turns out that one's inner life is configured in this albeit unknowable way. In general, luminosity failure makes for confusion or at least hesitancy in our normative theorizing. That said, we do not find the case as described an all things considered compelling one against the normative theory that we have described. After all, there is something to be said for the claim that the person who knows they have turned the coffee pot off should not be going back to check, that he ought not to treat his anxiety as a reason for going back. That is not to say of course that we would deem his going back inexcusable—but excusable lapses from a norm are no counterexample to that norm.

Objection 5. There are cases in which one has every reason to think one knows that *p*. In such a case, should not one treat the proposition that *p* as a reason for acting in a decision situation for which the truth or falsity of the proposition that *p* is relevant?

Reply. Here we see a strong analogy with the knowledge-assertion connection. Where one has every reason to think that one knows that *p* but does not, an assertion that *p* is quite excusable. But that is no objection to the norm that one ought to assert that *p* only if one knows that *p*. On the contrary, the need for an excuse in the case is explained by that norm. *Mutatis mutandis* in the action-theoretic setting.¹⁸

¹⁸ For a reply to this sort of objection against knowledge norms, see also p. 80 of Jonathan Sutton, *Without Justification* (Cambridge: MIT, 2007). Unger (*op. cit.*, pp. 209–10) gives another response to this sort of objection, directed against the claim that reasons for belief are known.

Objection 6. Consider someone who has a .4 credence that it will rain, and on those grounds brings an umbrella on their walk. According to the Reasons-Knowledge Principle, they are subject to criticism if they fail to know whether there is a chance that it will rain. But surely this person is rational, whether or not she acquires enough justification for her belief that there is a chance that it will rain. Perhaps she had to rush out the door, and did not have time to actually come to know what the epistemic chance of rain was. Surely, such a person is not subject to criticism.

Reply. There are certain responses to this kind of objection that are not available here. For example, Williamson, in defending the Knowledge Account of Assertion, considers the objection that someone may have to assert a proposition for reasons of expediency, despite the fact that they do not know that the proposition she asserts is true. Williamson (*op. cit.*, p. 256) responds by arguing that this is a case where a norm governing one kind of action, assertion, is “overridden by other norms not specific to assertion.” But this exact response is not available here. For the Action-Knowledge Principle is a norm governing action generally. So it is not an option that it can be trumped by norms governing action in some more general sense.

So, Williamson’s defense of the Knowledge Account of Assertion against a related charge is not available to us. Nevertheless, the criticism is not ultimately compelling. Someone might violate a norm, but nevertheless do so for an understandable reason. Here is an extreme case. Suppose I have inconsistent beliefs. It is understandable, if I am attacked by an axe-murderer, that I focus on defending myself against the axe-murderer rather than resolving the inconsistency in my beliefs. But the fact that expediency demands that I defend myself against the axe-murderer does not show that I do not violate a norm of rationality by having inconsistent beliefs. Similarly, the fact that expediency demands that I rush out of the door clutching an umbrella does not show that I do not violate a norm of rationality if I do not know that there is a chance of rain. The fact that our epistemic concepts play a central role in the appraisal of action is strong evidence that the Action-Knowledge Principle is a norm of rational action; the fact that we do not blame someone forced into a quick decision is no evidence at all against it.

Objection 7. Your decision theoretic framework strongly suggests that we do not know anything. After all, it’s clear that for any ordinary empirical claim, one should not bet on that claim at any odds and no matter how high the stakes. But insofar as I know that empirical claim, the decision theory will claim that I should bet on it at any odds and whatever the stakes. So, holding fixed the decision theory, it seems that one does not know of any empirical claim that it is true.

Reply. It bears emphasis that structurally similar concerns could be raised against myriad other decision theories. Suppose one is a subjective Bayesian. Consider now a case where one is subjectively certain that one feels cold but is not in a position to know that one feels cold. Suppose now one is offered a contract according to which ten children die if one does not feel cold but one gets a cent if one does. We can easily imagine an objector saying ‘One obviously ought not to accept the contract in that situation. Someone who did would be obviously susceptible to criticism. So either the decision theory is wrong or else we have to deny that we can be subjectively certain that we feel cold.’ Here, as with our own approach, it is hard to achieve a stable normative perspective owing to luminosity failures.

There is more that can be said, however, beyond general points about the confusing features of luminosity failure. We wish to outline two strategies for making our distinctive approach more palatable. We believe that one or perhaps both of these strategies will be helpful in diminishing the intuitive force of the original objection.

One strategy that we have defended in depth elsewhere is the *stakes-sensitive* strategy. If knowledge is constitutively related to one’s practical environment, then it is open to us to claim that while one may know that p in a situation where not much is at stake as concerns the proposition that p , one loses knowledge once one enters an environment where a good deal is at stake as regards the truth or falsity of the proposition that p . On such a view, knowledge—and hence probability 1—vanishes when the stakes go up. A decision theory of this sort will not make the claims about high stakes situations that were anticipated and deemed problematic by the original objector. Of course, this idea does not speak directly to versions of the original objection that were focused on heavily weighted *odds* as opposed to high stakes. Suppose the objector were to insist that one ought not to bet on any empirical claim at any odds (say three dollars as against a million millionth of a penny) and concludes that, holding the decision theory fixed, one ought to be a skeptic. Stakes-sensitivity considerations are of no use here. A proponent of our view who relied merely on the stakes-sensitive strategy will accept that one ought to accept such bets, appealing to general anti-luminosity considerations of the sort alluded to as a means to diffusing objectors.

A second strategy relies on distinguishing normative evaluation of a particular episode of practical reasoning from normative evaluation of epistemic character. The general shape of the distinction is familiar from broader normative theory—sometimes one does what is right but in doing so manifests traits of character that will foreseeably get one into trouble in none too distant scenarios. For any good case in which one

knows that p , there will be a range of subjectively similar possible bad cases where one does not know that p but for all one knows one does know that p .¹⁹ Suppose that the bad cases are ones that could fairly easily come about in the future, and are ones where accepting a certain p -dependent contract leads to bad consequences. Someone who accepts the contract in the good case will obviously be so constituted as to accept the contract in the bad case. This fact about his constitution—this trait of epistemic character—may be worrisome, especially to the extent that the consequences of contract acceptance are disastrous, and especially to the extent that bad cases are expected to arise. (Note that given anti-luminosity and overwhelmingly plausible empirical assumptions, pretty much any normative theory about what one ought to do will generate situations where one acts and reasons as one ought but manifests a constitutive tendency to not do as one ought in other situations. Someone who always bet the farm when he is subjectively certain will, given anti-luminosity and plausible empirical assumptions, be someone who will be disposed to bet the farm when he is not quite subjectively certain.) Our proposed strategy is to separate two kinds of normative evaluations that are easily run together—on the one hand, judgments about whether one acted and reasoned as one ought in a given situation and on the other, which defects of epistemic character were revealed by that situation. Even in acting and reasoning as one ought one can and often does manifest some such defect—notably a disposition to act incorrectly and even disastrously in a range of other possible cases. Note that even in a case where someone both knows that p and knows that they know that p , and accepts a contract on the basis of the truth of the proposition that p , they may reveal a constitution that disposed them to act inappropriately in none too distant cases.²⁰ The proponent of the second strategy claims that it is only by properly distinguishing these two dimensions of criticism that we can achieve a suitably stable and nuanced normative theory.

Overlaid upon suitable warnings about the confusing effects of anti-luminosity, it seems to us that one or perhaps some suitable admixture of these two strategies can keep the objector at bay, especially when one realizes the extent to which structurally similar objectors can easily be raised—with strong initial intuitive force—against pretty much any normative decision theory.

¹⁹ Similarly, there may be a range of counterpart cases where for some distinct proposition q , one believes that q by a relevantly similar process, does not know that q but for all one knows one does know that q .

²⁰ We leave it an open question whether such flaws in epistemic character always guarantee a failure of knowledge at some level of iteration in the KK hierarchy in the original case.

IV. CONCLUSION

Our ordinary conceptual scheme suggests a connection between reasons and knowledge that is altogether ignored by the standard decision theory and moral psychologists alike, who urge, in effect, that we reinvent ourselves with a knowledge-free approach to practical and theoretical policing. Such a radical revision is thoroughly unnecessary. The value of knowledge is due in part to its role as a norm for action; its status as such raises no concerns that are not already present with any other putative norm.

JOHN HAWTHORNE

University of Oxford

JASON STANLEY

Rutgers University