

Must we act only on what we know?
Clayton Littlejohn
cmlittlejohnATyahooDOTcom
(Draft of 8.4.2007)

What is the relation between knowledge and action? According to Hawthorne and Stanley, they are related as follows:

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 .

The claim that nothing beyond knowledge of p 's truth could be necessary for properly acting on p seems perfectly harmless. The notion of propriety we are concerned with is epistemic, not moral or prudential. What more could we possibly need to act on the belief that p ? Superknowledge? The reason that RKP is controversial is that it asserts that nothing short of knowledge could warrant acting on p . I shall first argue that it is a mistake to think that you cannot have warrant to act on p unless you know p and then try to show where the arguments offered in support of RKP go wrong.

In support of RKP, Hawthorne and Stanley write:

Consider ... 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.¹

It is hard to know what to make of this passage because of the qualification 'prima facie'. Let us ignore it for the time being and consider the proposal that:

Fault 1: In cases where you oughtn't Φ unless p is true, can be blamed for Φ -ing if you don't first know that p is true.

It is not hard to see how you might get from Fault 1 to RKP, but two problems arise if you combine these two theses. It is a messy affair trying to chart the relationships between justification and blame. Perhaps the least controversial thing we might about justification and blame is this:

Fault 2: If you can be properly blamed for believing p , you aren't justified in believing p .

It follows from Fault 1 and Fault 2 that if you fail to know p for any reason, you cannot be justified in believing p . Two troubling consequences follow.

Here is the first. According to the JTB analysis of knowledge, if Audrey is justified in believing p and her belief about p is true, she knows p . We all know that this analysis won't do. Audrey and Cooper are having lunch in the land of fake dollar bills. They drove into the land of fake dollar bills this afternoon. Neither knows that they are in the land of fakes. That is why we can say that they are justified in believing that they have cash and not counterfeit bills in their pockets. Audrey recalls that she owes Coop ten dollars, so she reaches into her pocket, pulls out ten dollars, hands it to Coop, and says 'Now we are even'. While her belief that she has repaid her debt is true and she is justified in her belief, she does not know that her debt has been repaid. Or, so the story goes. If, however, she does not

¹ Hawthorne and Stanley, "Knowledge and Action", pp. 2.

know that her debt has been repaid, it follows from Fault 1 that she can be blamed for acting on her belief that p and for believing p . In turn, it follows from Fault 2 that her belief that p cannot be justified. In turn, it follows that Gettier cases are not possible and the JTB analysis of knowledge was right all along.

Here is the second troubling consequence. It is unclear whether justification supervenes on our non-factive mental states. Those who think that epistemic justification does supervene on our non-factive mental states will typically also assert that the conditions that determine culpability and blameworthiness supervene on these internal conditions. That is to say, they will deny that it is possible for two people to be in the same non-factive mental states, both Φ , but only one can properly be blamed for having Φ 'd. Externalists about epistemic justification often accuse internalists of making the mistake of conflating this perfectly harmless claim about blame and the perfectly false claim that justification supervenes on the same internal conditions. As part of their error theory about how the internalists are mistaken about epistemic justification, they will say that the conditions that determine culpability and blameworthiness that do supervene on the internal states are distinct from the conditions that determine deontic status. The internalist's mistake about justification is due to their mistaken view that conditions you cannot be culpable for failing to take account of cannot affect the justificatory status of your beliefs. Justification and permissibility, they will say, can come apart from culpability.²

Suppose that the conditions that determine blameworthiness and culpability do supervene on the internal states. If you combine this supervenience thesis with Fault 1, you get the result that you can only blamelessly believe p if every possible internal duplicate of you knows p . This in turn commits you to an infallibilist conception of epistemic justification according to which it is permissible to believe p only if the reasons for which you believe entail p . Infallibilism about justification is quite clearly at odds with ordinary intuition.

At this point, I suppose Hawthorne and Stanley might say that I have ignored an important qualification. They did say that someone who acts on p without knowing p is, "prima facie negligent". If what they meant to say was that someone who violates RKP *appears* negligent without saying that they are truly negligent, then I do not see how blame judgments interact with knowledge at all. Their discussion of the interaction between knowledge and blame or judgments concerning negligence is just a distraction. What they should say (and seem to say in some passages) is that anyone who violates RKP by reasoning from p is failing to reason from the premises they are permitted to while adding that such failures might nevertheless be excusable. Let us assume that this is the picture that they are working with and take Fault 1 off the table.

If they reject Fault 1, they can avoid the two difficulties we have considered thus far, but rejecting Fault 1 won't save RKP. Consider this passage:

Consider also how knowledge interacts with conditional orders. Suppose a prison guard is ordered to shoot a prisoner

² If you say that the conditions that determine culpability supervene on the subject's internal states while denying that the conditions that determine permissibility supervene on the conditions that determine culpability, you might make some headway in dealing with the problems of moral luck. For discussion, see Michael Zimmerman, "Taking Luck Seriously", *The Journal of Philosophy* 99 (2002) pp., 553-76.

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.³

It is true that if a guard shoots a prisoner on a baseless hunch they can be faulted, but there is a world of difference between knowingly shooting a prisoner trying to escape and doing so on a baseless hunch. You do not need RKP to explain their observation. I take it that if RKP is right, then if the guard has good reason to believe mistakenly that a prisoner is trying to escape, the guard ought not shoot the prisoner. The fact that the guard was reasonable in assuming that they were doing what they ought is an excuse of the shooting, not a justification. This seems right. Suppose that as the guard raises the rifle to take a shot at Tobias. He looks just like a prisoner escaping. A second guard standing nearby the first knows that Tobias is really an aspiring actor spending the weekend in the prison preparing for his upcoming role in a film as frightened inmate number two. The second guard might mace the first guard to stop him from shooting Tobias. The reason it is not wrong for him to mace the guard knowing how painful it will be for the guard to be sprayed with mace is that the first guard is about to do something he ought not. He's lost the right to non-interference as a result. While RKP gets this sort of case right it seems it gets them right for the wrong reasons.

To see this, forget about cases of reasonable but mistaken beliefs. Forget about the cases of aspiring actors that look like prisoners trying to escape and think about aspiring escapees who surround themselves with aspiring actors. According to the order, George, who is a prisoner and not an actor, ought to be shot if he tries to escape. According to the order, Tobias, who is an actor and not a prisoner, ought not be shot. Assume George tries to escape. According to the order, the guard ought to shoot him before he makes his escape. Because unbeknownst to the guard there are aspiring actors like Tobias dressed like prisoners, the guard does not know George ought to be shot. He merely reasonably and correctly believes George ought to be shot. According to RKP, it is wrong to act on the one premise that could justify shooting George (i.e., that he is a prisoner trying to escape). George ought not be shot. It seems to follow that the guard ought to shoot George and ought not shoot George. That seems like a contradiction.

It is not hard to see that you are courting disaster if you combine RKP with the view that it is possible for there to be positive duties to Φ if p is true. If it is possible for circumstances to arise in which p is true but can't known to be true, it follows that you both ought to Φ and ought not Φ . There might be some way to sort out this mess available to the defenders of RKP, but it seems the easiest way to sort out this mess is just to deny RKP.

Earlier I suggested that if you combined RKP with Fault 1, you had to deny that Gettier cases were possible. That seems pretty costly. You do not have to pay that cost if you deny Fault 1. You cannot make all the problems you face with Gettier cases go away simply by denying Fault 1. It seems that if RKP is true, if someone fails to know p for any

³ Hawthorne and Stanley, pp. 2.

reason at all, they ought not act on the belief that p . Now, if your belief about p is mistaken, I can see that the consequences of acting on the mistaken belief might be so bad that we are inclined to say that your having acted on p is possibly excusable because of your non-culpable ignorance, but wrong nevertheless. If your belief about p is unreasonably held, I can see how your acting on p might manifest the kinds of bad motives or intentions that show that you can be faulted for having acted on p . It is not hard to grasp the normative significance of negligence or recklessness. What is hard to understand is the suggestion that there are considerations beyond those that have to do with the accuracy of your beliefs and the reasonableness of holding those beliefs that have an additional kind of normative significance. If Audrey hands Coop the ten dollars she owes him, her bills are genuine, and she has no reason to think anything is amiss, precisely what is it that was wrong with her acting from the belief that by handing that bill over she'd repay her debt? I cannot fathom it. From Coop's point of view, it is not as if he would care whether she repaid the debt knowingly at home or unknowingly in the land of fake bills. If I imagine myself as an outside observer who knows that Audrey does not know she'll repay the debt *merely* because she is trying to repay that debt in the land of fake bills, I am not at all inclined to think that the advisory judgment 'You ought not act from the assumption that you will repay that debt' is correct. It seems that Gettier cases not only show that knowledge is not a mere matter of justified, true belief but also that knowledge is more than needed for permissibly acting from a belief or using a belief in deliberation. Maybe JTB is also more than needed for properly acting from a belief, but it is surely sufficient.

We have seen reasons to think RKP must be wrong, which are reasons to be very suspicious of arguments that suggests otherwise. The first argument we are offered draws heavily on ordinary usage. Hawthorne and Stanley write:

Suppose ... 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 didn't know that the restaurant was here".⁴

It is natural enough for Sarah to say this and for us to construe this as criticism of Hannah. The case provides little support for RKP however, because Hannah's belief fails to constitute knowledge for a variety of reasons (e.g., her belief is really no better than a hunch and her hunch is mistaken). To test RKP properly, it seems we should consider three variants on the example:

Restaurant 1

Hannah and Sarah are trying to find a restaurant, at which they have time-limited reservations. Instead of asking someone for directions, Hannah relies on her usually impeccable memory and decides to go left. She has been eating at this restaurant regularly for years. After walking for some amount of time, it becomes quite clear that they went down the wrong street. Unbeknownst to Hannah, the

⁴ Hawthorne and Stanley, pp. 1.

restaurant had caught fire three days ago and was working from an alternative location two blocks away.

Restaurant 2

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. They find the restaurant just in time when Hannah declares, “That was lucky, I was just guessing that it would be this way.”

Restaurant 3

Hannah and Sarah are trying to find a restaurant, at which they have time-limited reservations. Instead of asking someone for directions, Instead of asking someone for directions, Hannah relies on her usually impeccable memory and decides to go left. She has been eating at this restaurant regularly for years. After walking for some amount of time, it becomes quite clear that they went down the wrong street. Unbeknownst to Hannah, owners of a rival restaurant managed to trick all the local papers and news outfits into running a story according to which the restaurant burnt down and would be serving at an alternative location.

In Restaurant 1-3, Hannah does not know that the restaurant is to the left. However, it is only in Restaurant 1 and 2 that it seems natural for Sarah to say, “You should not have gone down this street, since you did not know that the restaurant was here”. The defender of RKP cannot say that the reason it seems unnatural to say this in Restaurant 3 is that Hannah is blameless in that example, because she is blameless in Restaurant 1 and yet Sarah’s remark seems natural. The defender of RKP cannot say that the reason it seems unnatural to say this in Restaurant 3 is that there is no reason for Sarah to say this so long as they arrived at the restaurant because it is natural for Sarah to say this in Restaurant 2. I cannot see how either the original restaurant case or these modified versions provide any more support to RKP than they do the thesis that knowledge is merely a matter of true beliefs that are not baseless.

There is a perfectly reasonable explanation for this pattern that does not assume RKP. We often use ‘knows’ loosely as if it meant something like ‘true belief’ or ‘tenaciously held true belief’.⁵ Unfaithful lovers will speak this way when they think someone ‘knows’ of their secret rendezvous. Thieves speak this way of cops who ‘know’ about the heist they have planned. In conversational contexts like this, the propriety of using ‘knows’ does not depend upon what is known. That Sarah’s remarks only seem proper in Restaurant 1 and 2 suggests that we are dealing with conversational contexts like this. Observing how ‘knows’ functions in such contexts provides no real support for RKP.

The second argument offered in support of RKP involves lottery propositions:

⁵ For discussion, see Alvin Goldman, *Pathways to Knowledge: Private and Public* (New York: Oxford University Press, 2002), pp. 183.

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 ... The most natural explanation for why one can't act on these beliefs is that these beliefs are not knowledge.⁶

Are there alternative explanations for why one ought not act on these beliefs? Yes, but to be fair, they seem contrived. Still, I think a contrived explanation is called for because it seems RKP does not give us the explanation we are looking for.

Maybe it is controversial, but let us suppose you do not know you will lose the lottery.⁷ Consider two possible explanations as to why this is. First, you might think that the kinds of grounds you have for lottery beliefs do not put you in a position to know that the lottery belief is true. Second, you might think that lottery beliefs are unsafe. Because they easily could have been mistaken, they do not constitute knowledge. I would like to suggest that the safety-based explanation cannot be right. If that is right then while the proper explanation is consistent with RKP, it cannot be that RKP is needed to explain the fact that you ought not reason from the belief that you will lose the lottery.

There are two potential difficulties with the safety-based explanations. First, while beliefs in lottery propositions fail to satisfy safety, it is not obviously true that knowledge requires safety. If that is right and RKP assures us that knowledge is sufficient for using p in practical reasoning, you can sometimes reason from unsafe beliefs. Let us assume, however, that knowledge does require safety. By RKP, you ought not reason from unsafe beliefs. That does not seem right, however. Sometimes there are unsafe beliefs you ought not reason from. There are some unsafe beliefs it is not wrong to rely on in reasoning.

Let us distinguish between two kinds of lottery belief. An overt lottery belief is a belief that you cannot entertain without grasping that its truth depends upon the outcome of a lottery that you have purely statistical grounds for accepting. A covert lottery belief is a belief whose truth depends upon the outcome of a lottery you have no special information concerning but you are not in a position to appreciate that there is such a lottery that threatens to render your belief false. Neither overt nor covert lottery beliefs are safe, so if we combine the safety-based explanation and RKP, it follows that you ought to rely on neither type of lottery belief in reasoning. Because RKP seems to deliver the verdict that you ought not reason from covert lottery beliefs, RKP delivers the wrong verdicts and the wrong explanation for what is wrong with reasoning from overt lottery beliefs.

It is possible that you believe p where p is a covert lottery belief where I know the outcome of the lottery and know that the only reason your belief does not constitute knowledge was that it was 'threatened' by a lottery that turned out not to render your belief false. For example, suppose you wanted to surprise Hannah and Sarah and knew where they were planning on having dinner. What you do not know and I do know is this: Hannah

⁶ Hawthorne and Stanley, pp. 2.

⁷ We need to assume this if RKP has even a chance at explaining what is wrong with the reasoning above.

bought a lottery ticket. If you thought about it, you would realize that if Hannah won the lottery, she would go somewhere really nice to celebrate. So, your belief that she will be at the restaurant she has reservations at is unsafe and not one you know to be true. However, I know that her ticket lost and know that Hannah knows this, too. I am not inclined to say you ought not reason from your covert lottery belief. However, according to RKP, if I know you do not know p , whether you know you do not know p or not, I know that you ought not reason from p . That seems like it is not true when the belief that p is a true belief and a covert lottery belief.

If there are lottery beliefs that are not known but not wrong to reason from, then RKP must be wrong. That suggests that if it is true that you ought not reason from overt lottery beliefs, it is not because such beliefs do not amount to knowledge *per se*, but because of the kind of grounds you have available for such beliefs. While explanations that do not assume RKP might look contrived, the inadequacy of RKP might make a virtue of contrivance.

We have seen that there are a number of good reasons to think RKP must be wrong and that the arguments for RKP are far from compelling. Appeals to ordinary usage support nothing stronger than a JTB account of when it is proper to rely on a belief in reasoning. While I have not solved the lottery paradoxes, I have suggested that their solution does not assume RKP, in which case the desire to explain the perversity of reasoning from overt lottery beliefs should not lead us to say that you cannot be warranted in reasoning from that which you do not know.