

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

The Reason-Knowledge Principle (RKP)

Where your choice is p -dependent, it is appropriate to treat the proposition that p as a reason for acting iff you know that p .¹

To some, the claim that nothing beyond knowledge of p 's truth could be necessary for properly acting on p might seem perfectly harmless.² 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? Perhaps the main reason that RKP is controversial is that it asserts that nothing short of knowledge could warrant acting on p . Here, I shall argue that it is a mistake to think that you cannot have warrant to act on p unless you know p and address the arguments offered in support of RKP.

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 (572).³

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

¹ Hawthorne and Stanley, "Knowledge and Action," *The Journal of Philosophy*, CV, 10 (2008): 571-90. A choice between options is p -dependent iff the most preferable option conditional on the proposition that p is not the most preferable option conditional on the propositions that $\sim p$. In what follows, unadorned page references will be to this article.

² Jessica Brown ("Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning," *Nous*, XLII, 2 (2008): 167-89) criticizes the claim that knowledge that p is true is sufficient for properly treating p as a reason for action. Ram Neta ("Treating Something as a Reason for Action," *Nous* (forthcoming)) addresses Brown's criticism of RKP and defends the sufficiency claim while attacking the necessity claim. For a criticism of Neta's alternative to RKP, see Clayton Littlejohn, "On Treating Something as a Reason for Action," *Journal of Ethics & Social Philosophy*, www.jesp.org (February 2009).

³ Timothy Williamson suggests there might be a similar connection between ignorance and negligence in his *Knowledge and its Limits* (New York: Oxford, 2001), pp. 241.

Fault₁: In cases where you ought not Φ unless p is true, you can be blamed for Φ -ing if you do not first know that p is true.

Given the plausible assumption that it is not proper to treat p as a reason to Φ when you can be blamed for treating p as a reason to Φ , it seems Fault₁ does lend support to RKP. Note that judgments of blame, negligence, and the like also seem to interact with ascriptions of justification:

Fault₂: If you can be properly blamed for believing p , you are not justified in believing p .⁴

To deny Fault₂, you would have to say that the facts in light of which someone can be properly blamed for believing p do not threaten the justificatory status of that belief. This would be an odd stance to take for someone who argues for RKP by means of the assumption that if you can be blamed for treating p as a reason for belief, it is not permissible to treat p as a reason for belief. Problems arise for any view that incorporates both Fault₁ and Fault₂. Combined, these assumptions entail that if your belief that p is true is practically relevant, your belief cannot be justified unless it constitutes knowledge.⁵ To use their example, suppose that you should not use a needle unless it is clean. From Fault₁, it follows that you can be blamed for using the needle if you use it but do not know that it is clean. It seems, intuitively, that you can be properly blamed for using the needle only if you are not blameless in the belief that it is clean. It follows from Fault₂, that you cannot be justified in believing that the needle is clean. Two objections should suffice to show that we should not accept Fault₁ and Fault₂. Given that Fault₂ is relatively uncontroversial it seems that the objections below, if sound, give us good reason to deny Fault₁.

First, according to the JTB analysis of knowledge, if Audrey is justified in believing p and her belief is true, she knows p . We all know that this analysis will not do. Audrey and Cooper are on a cross-country trip and stopped this afternoon to have lunch in the land of fake dollar bills. 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. She reaches into her pocket, pulls out ten dollars, hands it to Coop, and says that they are now even. While her belief that her debt is repaid is true and she is justified in that belief, she does not know that her debt has been repaid. Or, so the story goes. If, however, she does not

⁴ I shall assume that a belief is justified iff it might be justifiably relied upon in deliberation, an assumption that seems acceptable to Hawthorne and Stanley.

⁵ Let us say that your belief that p is true is practically relevant only if there is some p -dependent choice

know that her debt has been repaid, it follows from Fault_1 that she can be blamed for acting on her belief that p . In turn, she can be blamed 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. The JTB analysis of knowledge must have been right all along.

Here is a 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.⁶ That is to say, they will deny that it is possible for situations to arise in which two subjects in precisely the same non-factive mental states Φ and only one of these subjects is properly blamed for Φ -ing. 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.⁷ 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.⁸

Suppose, as many do, that the conditions that determine blameworthiness and culpability do supervene upon a subject's non-factive mental 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 . Such a view about justification is quite clearly at odds with ordinary intuition.

⁶ See Alexander Bird, "Justified Judging," *Philosophy and Phenomenological Research*, LXXIV, 1 (2007): 81-110.

⁷ For a recent example of this, see Jonathan Sutton, "Stick To What You Know," *Nous*, XXXIX, 3 (2005): 359-96.

⁸ 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", *Journal of Philosophy*, XCIX, 11 (2002): 553-76.

Such a view, arguably, leads to skepticism given the plausible additional assumption that next to nothing we believe about the external world we believe on the basis of infallible grounds.

At this point, I suppose Hawthorne and Stanley might remind us that I have ignored an important qualification. They said 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 but they refrain from saying that the subject is truly negligent, I fail to see how blame judgments interact with knowledge. Their discussion of the interaction between knowledge, blame, and negligence would be just a distraction. Perhaps what they should say (and seem to say in some passages) is that anyone who violates RKP fails to reason in the way that they ought to. However, they might add, someone might not be blameworthy for having reasoned in a certain way if they are non-culpably ignorant of the conditions in light of which they fail to reason as they ought to. Let us assume that this is the picture that they are working with and take Fault₁ off the table.

If they reject Fault₁, they can avoid the two difficulties we have considered thus far, but rejecting Fault₁ will not save RKP. Consider this passage:

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 (572).

It is true that if a guard shoots a prisoner on a baseless hunch they can be faulted for doing this, 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 make the first guard to stop him from shooting Tobias. The reason it is not wrong for him to make 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 but not a prisoner, ought not to 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 to be shot. It seems to follow that the guard ought to shoot George and ought not to shoot George. That seems like a contradiction.

You are going to run into trouble if you combine RKP with the view that it is possible for there to be positive duties to Φ if p is true if p is the sort of thing someone can non-culpably fail to know to be true when it is. If it is possible for circumstances to arise in which p is true where p cannot be known to be true, it follows that you both ought to Φ and ought not Φ . That seems like a contradiction. Similar difficulties arise if you combine RKP with the knowledge account of assertion. According to the knowledge account, you ought not to assert p unless you know p . It follows that you have a conclusive reason to refrain from asserting p if you do not know p . It follows from this and RKP that you ought not to assert p unless your belief that you know p constitutes knowledge. Since not everything you know is something you are in a position to know that you know, a problem arises, which is that knowledge of p 's truth is not invariably going to ensure that you have sufficient warrant for asserting p . However, the view that knowledge of p 's truth is sufficient for having epistemic warrant for asserting p is surely more plausible than the view that knowledge of p 's truth is necessary for having that warrant. So, it seems you ought not to accept

both RKP and the knowledge account of assertion.⁹ It seems the easiest way to sort out these messes is to deny RKP.¹⁰

Earlier I suggested that if you combined RKP with Fault₁, you had to deny that Gettier cases were possible. That seems pretty costly. You can avoid paying that cost if you deny Fault₁, but by denying Fault₁ you do not avoid all the difficulties caused by Gettier cases. It seems that if RKP is true, there is a *prima facie* reason for anyone who fails to know *p* to refrain from reasoning from *p*. Now, if your belief about *p* is mistaken, I can see that the consequences of acting on the mistaken belief might be terrible. Because of this, we might be inclined to say that your having acted on *p* was wrongful. 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 see the normative significance of negligence or recklessness. It is not hard to see why someone might think that RKP rightly says that you ought not act from unreasonably held beliefs or mistaken beliefs.

What is hard to grasp is the idea 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 will 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 should not act from the assumption that you will repay that debt' is correct. It

⁹ Hawthorne (*Knowledge and Lotteries* (New York: Oxford, 2004)) and Stanley (*Knowledge and Practical Interests* (New York: Oxford, 2005)) both defend the knowledge account of assertion.

¹⁰ Someone might say that the contradiction can be avoided if we read RKP as saying that there is a *prima facie* duty to refrain from acting on what you do not know. There is no contradiction in saying you have *prima facie* duty to Φ and to refrain from Φ -ing. This move is *ad hoc*, however, unless we can specify some potentially overriding duty that would justify acting on beliefs not known to be true. I cannot see what such a duty might be since the justification for acting on beliefs not known to be true would have to be epistemic. Be that as it may, even if we weaken RKP so that it says that it is merely *prima facie* wrong to treat *p* as a reason for acting if you do not know *p*, the view still faces one final objection. The objection is, I believe, decisive.

seems that the very same examples that show that we cannot identify knowledge with true beliefs we are justified in holding show that knowledge of p 's truth is not needed to properly rely on p in practical reasoning. Surely we have all we need to rightly reason from p if our belief about p is true and not unreasonably held.

We have seen reasons to think RKP must be wrong, so we have good reason to be suspicious of arguments for RKP. 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" (571).

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:

Restaurant₁

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

Restaurant₂

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.”

Restaurant₃

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 Restaurant₁₋₃, Hannah does not know that the restaurant is to the left. However, it is only in Restaurant₁ and Restaurant₂ 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 Restaurant₃ is that Hannah is blameless in that example, because she is blameless in Restaurant₁ and yet Sarah’s remark seems natural. The defender of RKP cannot say that the reason it seems unnatural to say this in Restaurant₃ 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 Restaurant₂. 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 'firmly held true belief'.¹¹ 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 Restaurant₁ and Restaurant₂ suggests that we are dealing with conversational contexts like this. Observing how 'knows' functions in such contexts provides no real support for RKP. At least, it seems that observing how 'knows' functions in such contexts provides no more support for RKP than it does for the hopeless view that knowledge is merely a matter of firmly held true belief.

The second argument offered in support of RKP is contained in this passage:

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 (572).

It is true that RKP provides a natural explanation as to why this reasoning is perverse and that counts in favor of it. However, we can undercut this argument if we can show that RKP delivers the wrong verdicts in some lottery cases.

Suppose you do not know you will lose the lottery.¹² This seems to be a consequence of two things. First, knowledge requires safety. Second, lottery beliefs are unsafe. RKP seems to rely on a safety-based explanation as to why you should not reason from the premise that you will

¹¹ For discussion, see Alvin Goldman, *Pathways to Knowledge: Private and Public* (New York: Oxford, 2002), pp. 183.

¹² We need to assume this if RKP has even a chance at explaining what is wrong with the reasoning above.

lose a lottery. This explanation is one we ought to reject on the grounds that it is not always wrong to reason from unsafe beliefs.

Let us distinguish between two kinds of lottery belief. An overt lottery belief is a belief you have purely statistical grounds for accepting so that you cannot entertain without grasping that its truth depends upon the outcome of a lottery. A covert lottery belief is a belief whose truth depends upon the outcome of a lottery but does so in such a way that the believer is not in a position to appreciate this. Neither overt nor covert lottery beliefs are safe, so if we combine the safety-based explanation with RKP, it follows that you ought to rely on neither type of lottery belief in reasoning. However, it seems that it is not always wrong to rely on covert lottery beliefs in practical deliberation.

Suppose you wanted to find Hannah and Sarah and thought you knew where they were planning on having dinner. What you do not know and I do know is that Hannah 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. Your belief that she will be at the restaurant to which you are now heading is unsafe. If safety is necessary for knowledge, you do not know where they will have dinner. I know that her ticket lost. I know your belief does not amount to knowledge, but I know that it is true and not unreasonably held to be true. I am not inclined to say you should 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 should be in a position to know that you ought not reason from p . Because there are lottery beliefs that are not wrong to reason from, RKP sometimes delivers the wrong verdicts and never delivers the right explanation. We might not know why you should not reason from overt lottery beliefs, but we know better than to think that we need to appeal to RKP to explain why you should not reason from such beliefs.

We have seen that there are a number of good reasons to think RKP must be wrong and no good reason to think RKP must be right. The appeals to ordinary usage and observations concerning lottery beliefs do not seem to support anything quite as strong as RKP. The strongest view that seems to fit with the available data is:

The Reason-Justified True Belief Principle:

Where one's choice is p -dependent, it is appropriate to treat the proposition that p as a reason for acting iff you are justified in believing p and p is true.

This principle does not imply falsely that every lottery belief is a belief you ought not reason from and does not imply falsely that you ought never to deliberate from beliefs that fail to constitute knowledge for purely Gettierish reasons. Obviously, a full defense of such a principle is best saved for another time, but allow me to say this on its behalf. There is a long tradition according to which there are always two ways to do the wrong thing. You can go wrong by bringing about outcomes that ought to be avoided or by carrying on as if you do not have proper concern as to whether you bring about states of affairs we ought to avoid. Obviously, if we are evaluating beliefs on epistemic grounds, we are concerned with their accuracy. We are also concerned with the agent's concern for accuracy, and so it makes sense to say that if p misrepresents how things are or the subject arrives at the belief that p in such a way that shows she has no concern for the truth, it follows that the subject's belief about p is not proper and not the proper basis for further deliberation. RJTBP captures this idea. To go beyond this and suggest that your beliefs must not only be accurate and the result of careful reasoning but also must be non-accidentally accurate is a view that is hard to fathom, in part, because it is hard to see why reasons for belief or action demand anything more than proper respect and successful conformity. It is because RKP implies that reasons for belief and action demand more than this that the principle goes too far.