

ACTIONABLE INTELLIGENCE

When is it permissible to treat p as a premise in practical deliberation? Some say that knowledge of p 's truth is necessary and sufficient for relying on the belief that p in practical deliberation. Others say that there is nothing more and nothing less to permissibly relying on the belief that p in deliberation than being justified in believing p . Both accounts are wrong. Seeing why neither account is right leads us to the proper account of 'actionable intelligence'. If your choice is p -dependent, it is proper to treat p as a reason for acting iff your belief that p faultlessly and faithfully represents how things are.

INTRODUCTION

When it is appropriate to treat something you believe as a reason for action? According to one recent proposal:

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

The suggestion that nothing beyond the knowledge that p is true could be necessary for acting on p seems quite plausible. The relevant notion of propriety here is epistemic. We are interested in determining whether it is appropriate epistemically to rely on a belief in practical deliberation. What more than knowledge could be needed for that? I suppose the reason this proposal raises some eyebrows is that it says that if a belief falls short of knowledge for any reason, it oughtn't be used in practical deliberation and the proposition believed oughtn't be treated as a reason for action. I shall argue that RKP suffers from some serious difficulties. Its consequences are troubling and its motivation is far from compelling. If that is right, we should try to find something to put in its place.

I have a hunch. I suspect that most will say that we ought to replace RKP with something along these lines:

RJP: When your choice is p -dependent, it is appropriate for you to treat the proposition that p as a reason for acting iff you are justified in believing that p .

¹ John Hawthorne and Jason Stanley (Forthcoming).

This is not a baseless hunch. We all know people with internalist sympathies who will say that RKP repeats the mistakes of the knowledge account of assertion.² They will assure us that alternatives to RJP will be at odds with ordinary intuition. Given a standard assumption about what is involved in being justified in holding a belief, I shall argue that we ought to reject RJP. After making my pre-emptive strike against RJP, I shall offer a positive account of ‘actionable intelligence’ and say which beliefs are fit to provide us with premises for practical deliberation.

KNOWLEDGE IN ACTION

In this section, we shall look at the case for RKP and I shall offer my case against it. The first and most obvious argument someone might offer on behalf of RKP is as follows.

The Assertion Argument

- (1) Like action, assertion is governed by an epistemic norm.
- (2) It would be surprising if there were different epistemic requirements on properly acting on the belief that p and asserting that that p .³
- (3) It is appropriate to assert p iff you know p .
- (C) It is appropriate to treat p as a reason for action iff you know p .

In support of (1), it seems rather obvious that we criticize both assertions and actions on epistemic grounds. If someone asserts something without grounds, we criticize it. If someone acts first and thinks later, we think poorly of them. Let us assume (2) for the time being. I think we ought to reject (3) and deny that knowledge is the norm of assertion. Here is the problem with that response. It is, given (2), highly unlikely that someone who accepts this sort of argument would question (3). And, it seems that the case for (3) is reasonably strong. So, rather than simply say that the reason that we ought to reject this argument is that we ought to reject the knowledge account of assertion, I shall argue that if

² For a defense of the knowledge account of assertion, see Williamson (2000). For criticism, see Kvanvig (Forthcoming) and Lackey (Forthcoming).

³ Hawthorne and Stanley (Forthcoming: 19).

you accept the knowledge account of assertion, you ought to reject RKP. And that means that the assertion argument is about as bad an argument as one might offer RKP.

Suppose you know p but do not know that you know p . According to (3), it is not improper to assert p . However, according to RKP it is improper to assert p . Hence, (3) is incompatible with RKP. Why? According to (3), there is a conclusive reason to refrain from asserting p if you do not know p and permission to assert p if you know it. However, *that there is permission to assert p* is not something you are in a position to know, so RKP says that the belief that it is not wrong to assert p is not something that could serve as a premise in your practical reasoning. If you cannot act on the assumption that you are not wrong to act, your action is thereby improper. Thus, there is either *more* than knowledge of p 's truth is necessary for proper assertion or *less* than knowledge is necessary for proper action. Rather than argue from the knowledge account of assertion to RKP, we ought to argue from the knowledge account of assertion against RKP.

If the assertion argument does not work, there are others to consider. Hawthorne and Stanley note that, "blame [and] judgments of negligence ... 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".⁴ Perhaps they intend something along these lines:

The Argument from Fault

- (4) If S ought not Φ unless p is true, S can be faulted for Φ -ing if S Φ 's without knowing p .
- (5) If S can be faulted for Φ -ing, S's Φ -ing is *prima facie* wrongful.
- (6) If S's Φ -ing is *prima facie* wrongful on the condition that S does not know p and Φ 's anyway, properly acting on the belief that p requires knowing p .
- (C) When S ought not Φ unless p , S ought not Φ unless S knows p .

⁴ Hawthorne and Stanley (Forthcoming, 2).

If we ignore their qualification '*prima facie*' for the moment, it seems that (4) generates some significant problems. For I think it is plausible to think (5) is true, but the combination of (4) and (5) commits us to two things we ought to reject.

Suppose we say, as is surely plausible, that if you can be properly blamed for believing p , you are not justified in believing p . What (4) says is that if you do not know p , you can be blamed for having acted on the belief that p . It seems intuitive to say that beliefs that you ought not deliberate from are beliefs that are not justifiably held. (You must say this if you accept RKP.) But, that means that beliefs that do not constitute knowledge are beliefs you can be faulted for holding. In turn, it follows that Gettier cases are not possible.

If Gettier cases are counterexamples to the JTB analysis of knowledge, it is because it is possible for S to be justified in believing a true proposition without knowing that it is true. The assumptions about fault needed for (4) tell us that such a person believes a proposition that they ought not rely on in practical deliberation *because* they can be faulted for relying on this belief in deliberation. As beliefs such as this cannot be justified, Gettier cases are impossible. The JTB analysis would be vindicated. We ought to reject (4).

To be fair, I ignored a qualification. Hawthorne and Stanley said that someone who acted on the belief that p without knowing p was '*prima facie* negligent', not that they were negligent. But, we have to ask what is left of the argument if we take account of this qualification. If the way to save the argument is to say that the argument does not assume (4) it is entirely unclear how we are supposed to take the alleged link between knowledge and fault as grounds for RKP.

A third argument for RKP is found in this passage:

Consider ... 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 condition in order to discharge it.*

What are we to make of this? It is true that the guard acted wrongfully if he shot a prisoner on a baseless hunch. But, there is more to knowledge than a belief that is based on more than a baseless hunch and that also happens to be correct. So, while we can agree with the intuition that the shooting of the prisoner is wrongful, that fact does not need RKP for the proper explanation.

In fact, we can use cases of conditional orders to show that if we have reason to do anything, we have reason to reject RKP. Warden Jeffries has decided to let Tobias spend the weekend in jail so that he might observe the other prisoners and prepare for a role in an upcoming movie. He is an actor dressed as a prisoner, not a prisoner. Jeffries forgets to alert the guards. Tobias spends the weekends doing the sorts of things that other prisoners do. At one point, he acts just like a prisoner trying to escape. He is, however, not a prisoner trying to escape but an actor acting like such a prisoner. Meanwhile, George, a prisoner, is trying to escape. According to the orders, the guards ought to shoot George but ought not shoot Tobias. The guards do not know that George is trying to escape. (If you do not have that intuition, keep adding aspiring actors until we have something like a fake barn case). According to the order, the guards ought to shoot George. According to RKP, the guards only ought to shoot George if they know that he is a prisoner. So, it seems that given the order and given the principle, the guards ought to shoot George and ought not shoot George. According to RKP, it is wrong to act on the one premise that would allow shooting George. So, if there are any positive duties or any reasons to do things rather than merely refrain from doing things, among the things we have reason to do is reject RKP.⁵

⁵ I think that the combination of RKP and the claim that we have positive duties even when we cannot know that we have such duties commits us to a contradiction. It commits us to saying that on the occasion that there is a duty to Φ provided that C obtains when one cannot know C obtains, one ought to Φ and ought not Φ . Maybe a defender of RKP can say that there is no contradiction here because the 'ought' is not an all things considered ought. Perhaps they will say that *prima facie* you ought to Φ and *prima facie* you ought not Φ . This move seems *ad hoc*. It makes sense to posit *prima facie* duties or obligations when it makes sense to posit conflicting duties that can override each other or duties that have certain kinds of canceling conditions. The case I've described seems rather unlike the ordinary cases of conflict or canceling. Someone who saw that the order required the guard to shoot and that the guard did not know that the prisoner was trying to escape would likely not be inclined to say that there was a conflict here akin to the conflict between, say, fidelity and beneficence. Someone who saw that the order required the guard to shoot and that the guard did not know this would not likely see the duty to conform to RKP as canceled in the

Part of the case for RKP rests on observations about how we use ‘knows’ in criticizing the way in which people act on certain beliefs. We are asked to imagine that Hannah and Sarah are trying to find a restaurant and are under some pressure to make it in time for their reservation. If Hannah acts on a hunch and takes them down the wrong street, it is natural for Sarah to say ‘You shouldn’t have gone this way since you didn’t know this would take us to the restaurant’.⁶ In fact, if Hannah acts on a hunch and takes them the right way, Sarah can still complain. As I said earlier, there is far more to knowledge than getting things right by believing something that is more than a lucky guess. If it is really the fact that Hannah’s failure to know something entails that she oughtn’t rely on that as a premise in deliberation or act on it, we ought to be able to elicit the intuition that Hannah ought not act beliefs that fail to constitute knowledge for purely Gettierish reasons. Consider this example:

Sarah and Hannah are on a long road trip and are a bit pressed for time. They start driving in shifts. It is Sarah’s turn to drive and Hannah is taking a nap. Sarah is hungry and sees a diner from the highway, exits, pulls into the parking lot, and wakes Hannah. Sarah and Hannah order their lunches and as they eat their sandwiches, Hannah asks where they are. Sarah tells them that they’re in Ipswich. Hannah remembers that Ipswich is the land of fake diners. The hills of Ipswich are filled with very convincing fakes. Many a weary, hungry traveler has pulled over to grab a bite to discover that they had exited because they saw a fake diner. Hannah knows Sarah did not know that the diner was real. She was merely justified in believing correctly that the diner was real.

In such a case the remark, ‘You should not have pulled off at that exist because you did not know there was a diner here’ seems out of place. It surely is not the case that Sarah ought to have acted other than she did. If we assume that her only possible reason for pulling off

way that a friend might let you out of a promise. The argument from conditional orders seems like a strong argument against RKP.

⁶ Hawthorne and Stanley (Forthcoming, 1).

could have been that there was a place to get something to eat at the exit she took, it seems we have two things. First, we have a counterexample to RKP. She used as a premise *by taking this exit, we shall get something to eat*, but she did not know that this was true. Nevertheless, it seems it was not wrong for her to act on this premise. Second, we have good reason to think that there is something dodgy taking place in the examples Hawthorne and Stanley are using to motivate RKP. It seems that in the kinds of restaurant cases they describe, the remark to the effect of ‘If you didn’t know, you shouldn’t have done that’ seems appropriate only if the reason that the relevant proposition was not known was that the belief in it was baseless or the proposition is false. That you cannot properly say this sort of thing when the sole reason the proposition was not known was due to Gettierish reasons suggests that these are contexts in which ‘knows’ is being used in a peculiar way.

It seems no more appropriate to use cases to establish the link between ‘knows p ’ and ‘permissibly treats p as a premise’ in which the use of ‘knows p ’ is insensitive to considerations beyond the truth of p or the reasonableness of having assumed p than to characterize the extension of ‘knows p ’ using such cases. So, the intuitive description of these restaurant cases, which are cases in which the propriety of use of ‘knows p ’ cannot be challenged on Gettierish grounds, should no more serve as a basis for an account of the conceptual connection between knowledge and permissibly using something as a premise than it should serve as a basis for adopting the JTB account of knowledge. Unless we are going to give the JTB account serious reconsideration, I cannot see that these restaurant cases give us serious reason to think RKP is true.

Perhaps the strongest argument for RKP is this:

The Lottery Argument

- (1) It is absurd to use as a premise in practical reasoning lottery propositions (e.g, this ticket for tomorrow’s lottery will lose).
- (2) The most natural explanation for why one cannot act on these beliefs is that these beliefs are not knowledge.
- (C) Thus, we ought not act on a belief unless it constitutes knowledge.

Let us suppose that you cannot know lottery propositions without insider's information or after the drawing. And, let us suppose that (1) is true. It is true that a natural explanation of (1) is that the relevant proposition is not known. Perhaps it is also true that, *ceteris paribus*, a natural explanation is preferable to a contrived one. However, we have also seen independent reason to think that RKP is false. Myself, I prefer a contrived explanation containing true propositions to the natural ones containing false ones. But, suppose that a defender of RKP is not yet convinced that RKP is false. What then? I will argue that the RKP offers us the wrong explanation because even if we assume that no lottery propositions are known, the use of lottery propositions in practical deliberation is not wrongful in every case. So, if that is right, we can run a lottery argument against RKP on the grounds that it delivers the wrong verdicts in some lottery cases.

To illustrate (1), Hawthorne offered us this piece of reasoning:

[In the upcoming 10,000 ticket lottery with a \$5,000 prize in which I own only one ticket] 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.

Let us grant that this piece of reasoning is absurd. Consider a second piece of reasoning:

We both want a drink.

If we go to the bar, we will not make it for last call.

If we go back to my place, we have drinks in the fridge.

So, we ought to skip the bar and head to my place.

Suppose the CIA runs a lottery in which it selects at random houses to raid for beer. Suppose I know about the lottery and know you know nothing of it. Suppose I know you have a good memory so you would remember if you had stocked up recently. Now, suppose I know that your number did not come up and so the CIA did not raid your kitchen. I know that you reasoned from a lottery proposition without knowing it to be a lottery proposition. It seems that the belief that there would be beer at your place does not constitute knowledge for the same reason that typical lottery beliefs do not. Yet, it does not seem from my perspective there is anything wrong with your reasoning. If RKP treats both pieces of reasoning the same and yet only one seems wrongful, it cannot be that the reason the first piece of reasoning is wrongful has to do with considerations that are not salient to the

subject. Rather, the wrongness has to be due to something accessible to the agent. And that suggests that while RKP undoubtedly gets some of these cases right, it is doubtful that they get them right for the right reason.

We have seen a number of ways in which RKP is problematic. In each of these cases, the problems have to do with the way in which information inaccessible to the individual reasoning from a premise shows that the belief the individual reasons from does not constitute knowledge. Intuitively, it seems that when the considerations hidden from the individual prevent that individual from knowing p , it does not follow that these considerations prevent that individual from properly reasoning from p . There seems to be a natural fix. We can move to an account in which the permissibility of relying on p in practical deliberation turns on factors accessible to the subject. We can move to an account on which the permissibility of relying on p in practical deliberation turns on whether S is justified in believing p in virtue of how things seem to her.

JUSTIFICATION JUSTIFIES

What could justify relying on a belief in practical deliberation? Justification. That seems like a perfectly natural answer. If you distinguish someone's being justified in believing p from someone's knowing that p is true, the arguments from the previous section suggest that a justification based approach will be preferable to a knowledge based approach. In this section, I shall argue that given a natural assumption about justification, we can reject the following account:

RJP: When your choice is p -dependent, it is appropriate for you to treat the proposition that p as a reason for acting iff you are justified in believing that p .

It might seem that to evaluate RJP we have to have some sense of when someone is justified in believing p . We do not. To show that RJP is false, I will only have to assume something about what justification is not. I will have to assume that it is possible for you to be justified in believing p even if p is not true. I shall offer two cases show that regardless of what you take justification to be, if it is possible that you could be justified in believing something false, RJP cannot possibly be right.⁷

⁷ Not everyone grants that assumption. Sutton (2005) defends the view that S is justified in believing p only iff S knows p . On his view, RJP and RKP amount to the same thing.

Let's begin with a story:

Loan Shark

Moro is behind in his payments to Snake, the loan shark. Snake gave Moro a severe beating last week, with a warning that, if Moro missed another payment, he would be killed. The payment is due today but Moro doesn't have the money. He borrows a can of mace and hangs out at Deffi's Deli hoping Snake will leave him alone in public. He is shocked when he sees a man with a menacing look he takes to be Snake come in and walk straight towards him. Moro says, 'I won't let you get me Snake!' and he pulls out his can of mace and takes aim.⁸

Suppose the story continues as follows:

A

Deffi has a can of mace. He knows that the man Moro takes to be Snake is Snake's twin brother. He also knows that while his twin looks dangerous, he isn't a threat to anyone. Deffi knows that to disarm Moro, he'll have to mace him. He does so, intervening on behalf of Snake's brother.

Concerning this case, it seems natural to say:

- (1) Deffi is justified in intervening on behalf of the man Moro intends to Mace.
- (2) Moro has lost a right to non-interference he'd otherwise have if he didn't intend to mace the guy.

From (2), I think it makes sense to say:

- (3) Moro should have acted otherwise/Moro should not have acted on the belief that it was not wrong to mace the guy he saw.

But, there is nothing in the story that would compel us to say that Moro was not justified in believing that grounds of self-defense would permit him to mace the guy approaching, or

⁸ My story is taken from Robinson (1996). I cannot assume that he would approve of how I am using his story, but I suspect he would be sympathetic.

that he ought to mace the guy approaching, or that it would not be wrongful of him to mace the guy approaching. The only grounds we would have for saying such things is that his belief about the identity of the man approaching is mistaken. So, denying that he was justified in believing that he ought to mace the guy approaching would contradict our assumption. And, because Moro ought to have acted otherwise, he acted without sufficient justification. So, his acting on the belief that he had sufficient justification was acting without sufficient justification. And, so it seems that if we insist that he was justified in his belief, this is insufficient for relying on that belief in practical deliberation if we assume that there are rational requirements linking judgments about what one ought to do with intentions to do what one judges one ought to such as this:

- (4) If it is not the case that one ought not judge that one ought to Φ it is not the case that one ought not intend to Φ .

Someone might object that the argument assumes that Moro lost the right to non-interference because he was in the wrong. But, someone might say, while it is clear that Moro has lost the right to non-interference and Deffi was not wrong to intervene in a way that would predictably injure Moro, it is possible for someone to lose the right to non-interference without engaged in a course of action that they ought not have. To see this, just consider:

Air Raid

Cooper intends to bomb a munitions factory. Destroying the factory is an important step towards winning a just war. If the bombs hit the factory, the explosion will destroy the apartments that abut the factory. Audrey is among the non-combatants who live in those apartments whose deaths count as acceptable losses. Audrey realizes that she can defend herself and her family by manning an anti-aircraft gun and firing at Cooper's plane.

It seems that while a just war theory might allow Coop to drop the bombs even if the bombs would kill Audrey and her family, it would not forbid Audrey from taking up arms against Coop to defend her and her family so that they are not killed as side-effects. So, while Coop is not acting wrongfully and might have the right not to be interfered with by just anybody,

Audrey seems to have the right to interfere with his action even if that involves firing on his plane. So, we cannot say that if someone loses the right to non-interference that loss is due to the initiation of an impermissible course of action.

To get around this worry, just consider a second way of continuing *Loan Shark*:

B

Deffi has a can of mace. He knows that the man walking in is Snake and that Snake is here to kill Moro. He also knows that Moro's can of mace is empty, which means that Moro cannot defend himself. As Snake approaches Moro, Deffi maces Snake to protect him.

It seems that the range of permissible options open to Deffi are much more limited in *Loan Shark B* than the range of permissible options open to Audrey in *Air Raid*. It seems intuitively that Audrey is permitted to take arms to defend her family, to do nothing, or perhaps to take arms against someone who was taking arms against Coop thereby sacrificing herself and her only chance for defending herself for a just cause. In *Loan Shark B*, Deffi either must refrain or must take arms against Spike. In *Loan Shark A*, Deffi must either refrain or must take arms against Moro. The best explanation as to why there is this difference in the range of permissible options, I submit, is that in *Loan Shark A*, Moro's intended course of action was wrongful but it was not wrongful in *Loan Shark B*. And, if that is right, then the beliefs that Moro might permissibly act from differ in *Loan Shark A* and *Loan Shark B*. However, the beliefs he is justified in holding in *Loan Shark A* and *Loan Shark B* do not differ. So, while RJP says he has sufficient justification for using his belief that he is permitted to spray the guy approaching with mace in both cases, we know that this is not true.

The problem with RJP, I submit, is that it says that we are permitted to use beliefs that we are justified in holding on occasions where acting on those beliefs would amount to performing actions that are excusable wrongs. However, if they are excusable wrongs, the actions are impermissible. So, unless a defender of RJP wants to say that permissibly relying on *p* in reasoning is one thing and permissibly acting on the belief that *p* something else entirely, RJP turns out to be incompatible with our moral intuitions. Here is one more case to drive this point home.

Chemist

Chris, a chemist, has just mixed a concoction that is supposed to help with heartburn. Assuming that it is possible to have a justified, false belief that this concoction will help with heartburn without causing any unfortunate side effects, assume that Chris has such a belief. He gives the mixture to Alphonse. Meanwhile Bertrand mixes up a batch and takes it for his heartburn. Alphonse and Bertrand break out in a painful rash because they took this mixture. It turns out that the concoction does have an unfortunate side effect. Luckily, Chris knows a cure for the rash.

Intuitively it seems that if he has only enough ingredients to mix up one cure, Chris should give it to Alphonse rather than Bertrand. Wouldn't it be wrong for Chris to flip a coin to determine who gets the stuff? If we assume Alphonse has a stronger claim on Chris' assistance, we have another argument against RJP. According to RJP, Chris was justified in believing that it was not wrong to give Alphonse the concoction. However, in light of what happened, it is clear that *if* he helps Alphonse this is no mere act of beneficence. If it were merely an act of beneficence, it is unclear why there is this difference in the stringency of the duties he had concerning Alphonse and Bertrand. Yet, we are assuming that there is this difference in stringency. In order to explain this difference, it seems we have to say that Alphonse's giving Bertrand the stuff is a matter of discharging a reparative duty. Since reparative duties are responses to previous wrongs, we have to say that the initial action that Chris was justified in believing was not wrongful *was* wrongful. And that is to say that while he was justified in his belief it was not permissible to act on that belief or give that belief the role it had in practical deliberation.

In response to these sorts of cases, I suppose someone who defended RJP might say this. We have to grant that the conditions under which someone is justified in believing p ensure that the conditions are such that it is not wrong to act on the belief that p or treat p as a premise in practical reasoning. However, intuition tells us that just as the internal conditions suffice to make it the case that someone is justified in believing something, intuition tells us that internal conditions make it the case that someone is justified in acting

on a belief.⁹ Thus, the internal conditions make it the case that someone is permitted to use such internally justified beliefs as sufficient grounds for action.

I think two responses to this response are sufficient. First, there is a way of squaring internalist intuitions about when someone is justified in believing something (e.g., that it is permissible to Φ) with externalist intuitions about the conditions under which someone is permitted to Φ . We can deny RJP. Perhaps we can say that conditions for justification and conditions for permissibility differ. Or, perhaps we can say that a belief might satisfy the norms of belief but not be a belief that is properly included in deliberation. Rather than doing that, I want to challenge the appeal to the intuition that internal conditions suffice to justify our actions. I know that there are epistemologists who report these intuitions. I know that some epistemologists aim to offer a theory that accommodates folk intuition or the intuitions of ordinary folks. I suspect some think that their internalist theory of justified action would do both. If they think that, the evidence suggests that they are wrong. Some empirical research suggests that the ordinary or lay conception of justified action is an externalist one on which the justifiability of a course of action depends upon its objective characterization.¹⁰ So, while I am sure that some internalists who want to defend RJP will say that any proposed account of the beliefs we are permitted to act on or rely on in practical deliberation will run afoul of, say, the intuitions underlying the new evil demon problem, I think we ought to ask whose intuitions those are. If they are the intuitions of philosophers wedded to a theory, I'm not sure that these are the intuitions we ought to be accountable to. Not if these intuitions diverge from folk intuition, they force us to revise the intuitions

⁹ Lackey (Forthcoming), for example, appeals to the new evil demon intuitions to say that you can have sufficient warrant for asserting false propositions. Audi (2001) and Wedgwood (2002) both seem to think that an action can be justified in virtue of the internal grounds the agent has for performing it and that the justification for the action does not turn on the external or objective character of the action.

¹⁰ See Darley and Robinson (1998) for defense of the view that the folk notion of justified action is one on which the action having the status of being justified depends upon the external characterization of the action or its objective character. I know of no further empirical research into this question that would support the view that the internalist is right about what it takes for an action to be justified. For criticism of such internalism concerning *legally* justified action, in addition to Darley and Robinson (1998), see Moore (1993) and Gardner (1996). For a useful discussion of the external characterization of action and the question as to whether such a characterization affects whether it is permissible or excusable, see Zimmerman (2004).

reported above about the moral status of our actions, and they make hash of distinctions between the justified or permissible actions and the excusable wrongs.

WHAT IS ACTIONABLE INTELLIGENCE?

I believe that the proper characterization of actionable intelligence is as follows:

RJTP: When your choice is p -dependent, it is appropriate for you to treat the proposition that p as a reason for acting iff you are justified in believing p and p is true.

It seems that such an account gets the Gettier cases right. In such cases, the beliefs that are not known for purely Gettierish reasons are not invariably wrong to act on. It seems that such an account gets the lottery case right. We know that you oughtn't believe in lottery propositions and we know that this must be due to the fact that you are not justified in believing lottery propositions even if they are true. If we can say that it is unreasonable to believe p if you are not justified in believing p , the virtue of RJTP is that it helps itself to whatever RJP says about lottery propositions. We know that RJTP makes sense of the moral intuitions that RJP cannot. It seems there is something going for the view. Perhaps a piece of actionable intelligence is really a belief that is not wrong to act on or deliberate from, and perhaps the beliefs it is not wrong to deliberate from or act on are the beliefs that faithfully and faultlessly represent how things are. If that is right, RTP is the right analysis of actionable intelligence.

The view might seem cobbled together. There is something unsatisfying about a view that seems as if it is patched together from the bits and pieces left over from the views that we rejected. It might seem far less elegant compared to RJP or RKP. I actually think that RJTP follows from a simpler view. It follows from two things: the assumption that a belief is properly held and properly deployed in reasoning iff it satisfies the norms of theoretical reason; the assumption that the fundamental norm of theoretical reason is the truth norm (i.e., that you oughtn't believe p unless p is true).

How do we get from these assumptions to RJTP? The truth norm states a necessary condition on permissibly holding a belief. It only says that we oughtn't believe what isn't true. To see how we get from here to RJTP, note that we tend to think that whenever there is a rule, principle, or norm that places some requirement on us to refrain from Φ -ing unless C , our Φ -ing might be wrongful in one of two ways. Most obviously, we might Φ even if C

doesn't obtain. If that happens, Φ -ing is wrong in virtue of contravening the norm. It also seems that in virtue of this very requirement we might wrongfully Φ because we should not have expected to conform to that requirement's demand.¹¹ If there was no good reason to take that risk and given how things seemed to us, we had no reason to assume that C obtained, our Φ -ing shows a kind of failure to respect C's status as a requirement on Φ -ing.

That is all very abstract. Let me offer an example. Because of the kinds of things kids are, you oughtn't run them over. Suppose you're backing out of your driveway, your neighborhood is full of kids, but you don't stop to check to see that there are no kids behind you. On its face, it seems that the action is wrongful for the very same reason that you oughtn't run over kids. Or, if we don't have two aspects of the same reason, it seems we have two related reasons that are generated by the same norm. Or, if we don't have multiple reasons generated by the same norm, we have a subsidiary norm that is generated by the original norm. If you oughtn't Φ unless C, perhaps there is an additional norm in virtue of which you oughtn't Φ unless you can reasonably assume that C obtains.

If we say that for every fundamental norm of the form 'You oughtn't Φ unless C obtains', there will be reasons to refrain from Φ -ing without C obtaining and reasons to refrain from Φ -ing unless you can reasonably assume that C obtains, we can say that there are two kinds of wrongs grounded ultimately in the same thing. Some wrongs imply fault. Some wrongs imply the transgression of the fundamental norm. They are conceptually connected. Applied to the case of the truth norm, if a belief is false, it is wrongful in the second sense. If it is true but cannot be assumed to be true, it is wrongful in the first. According to RJTP, it is permissible to act on the belief that p if it is not wrongful in either way. If it is wrongful in either way, the acting on such a belief is *prima facie* wrongful.

¹¹ Williamson (2000: 246) says that you cannot consistently say that assertion is governed by the norm of truth and deny that assertion of p 's truth demands evidence for believing p for roughly this reason. For an excellent discussion of wrongs and faults, see Gardner (2005). For criticism of this sort of distinction and the derivation of an evidential norm from a truth norm, see Lackey (Forthcoming). Lackey criticizes the use of the distinction between 'primary' and 'secondary' propriety in Weiner's (2005) defense of the truth account of assertion. She argues that the distinction is spurious, but does so using examples drawn from sports. While I agree that there is no primary/secondary propriety distinction in basketball, I think that does not give us any reason to doubt the distinction since in creating rules for games like basketball, you are at liberty to decide whether the application of rules is guided by distinctions we would normally recognize in morality or feel compelled to recognize under the law.

If we assume what I think is a traditional view about wrongs, reasons, and norms we can state the argument against the knowledge account as follows. According to RKP, anytime you believe p and rely on p in practical deliberation without knowing p , your reasoning from p is wrongful. The tradition recognizes two types of wrong. Some wrongs are fault-implicating wrongs (Call these ‘wrongs-1’). Some wrongs involve transgressing a norm (Call these ‘wrongs-2’). If we say that the wrong of acting on p without knowing p is wrongful-1, we end up saying that any instance of ignorance is culpable. Among other things, we have to deny that Gettier cases are possible. If we say that it is wrongful-2, we have to say that failing to know for purely Gettierish reasons constitutes a conclusive reason against reasoning from p . But, that clashes with intuition. Moreover, we tend to say think of wrongs-2 as incorrect, or as mistakes, or as things that ought to be otherwise. But, this is not what we find. Unless the tradition failed to notice a kind of wrongfulness that is distinct from wrongfulness-1 and wrongfulness-2, we know there is something wrong with knowledge accounts across the board.

As for RJP, it seems to collapse the distinction between wrongs-1 and wrongs-2. Rather than say implausibly that external circumstances can have some bearing upon fault, people who go this route tend to go ‘internalist’. But then they cannot do justice to the kinds of intuitions we discussed and cannot say that the agents described above whose actions were wrongful really were wrongful. Conceptually it seems that their account does not allow for the possibility that external circumstances can have a bearing upon the permissibility of an action. Among other things, it seems they end up misclassifying excusable wrongs as permissible actions. If they are to avoid doing that, they have to give up on the entire project of connecting the conditions under which a belief is permissibly held or the conditions under which someone is justified in holding a belief to the conditions under which someone can permissibly act on a belief. Having given up on that project, they have to abandon RJP.

REFERENCES

- Audi, R. 2002. “An Internalist Theory of Normative Grounds.” *Philosophical Topics* 1&2: 19-45.
- Darley, J. and P. Robinson. 1998. “Testing Competing Theories of Justification.” *North Carolina Law Review* 76: 1095-1143.

- Gardner, J. 1996. "Justifications and Reasons." In A. Smith and A. Simester (eds.) *Harm and Culpability*. New York: Oxford University Press.
- _____. 2005. "Wrongs and Faults." In A. Simester (ed.), *Appraising Strict Liability*. New York: Oxford University Press.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. New York: Oxford University Press.
- Hawthorne, J. and J. Stanley. Forthcoming. "Knowledge and Action." *Journal of Philosophy*.
- Kvanvig, J. Forthcoming. "Knowledge, Assertion, and Lotteries." In *Williamson and His Critics*. New York: Oxford University Press.
- Lackey, J. Forthcoming. "Norms of Assertion." *Nous*.
- Moore, M. 1993. *Placing Blame*. New York: Oxford University Press.
- Robinson, P. 1996. "Competing Theories of Justification: Deeds vs. Reasons." In A. Smith and A. Simester (eds.) *Harm and Culpability*. New York: Oxford University Press
- Sutton, J. 2005. "Stick to What You Know." *Nous* 39: 359-96.
- Wedgwood, R. 2002. "Internalism Explained." *Philosophy and Phenomenological Research* 65: 349-69.
- Weiner, M. 2005. "Must We Know What We Say?" *The Philosophical Review* 114: 227-51.
- Williamson, T. 2000 *Knowledge and its Limits*. New York: Oxford University Press.
- Zimmerman, M. 2004. "Another Plea for Excuses." *American Philosophical Quarterly* 41: 259-66.