

INTRODUCTION

We often speak as if the concept of knowledge plays a central role in normative appraisal. Whether you should tell someone that the bridge is safe to cross or treat the (apparent) fact that the stuff in the bottle is gin as a reason to serve it to a friend seems to depend upon what you know. If you see the ropes snap and the bridge give way, you learn too late that you shouldn't have told them it was safe to cross. If you serve your friend petrol and tonic thinking the stuff in the bottle was gin, you will soon discover that you shouldn't have served that drink. As a rule, what you can properly treat as a reason for action or tell others depends upon what you know:

KPR: Whether it is proper to treat p as a reason for action depends upon whether you know p .¹

KA: Whether it is proper to tell others that p depends upon whether you know p .²

From a certain theoretical perspective it's hard to see how knowledge could play this role in normative appraisal. Doesn't justification play this role? If it's possible to justifiably believe what you don't know, KA implies that you shouldn't tell others what you rightly believe. KPR implies you that your justified beliefs are not the justified basis for action. Many epistemologists balk at this. They insist that justification, not knowledge, is the norm for practical reason and assertion:

JPR: Whether it is proper to treat p as a reason for action depends upon whether you justifiably believe p .³

JA: Whether it is proper to tell others that p depends upon whether you know p .⁴

Those who defend KA and KPR often complain that justification norms like JPR or JA are too weak to account for various intuitions.⁵ Among the cases thought to cause trouble for JPR and JA are cases of reasonably believed falsehoods. If we assume that one of the orthodox accounts of justification is correct, JA and JPR imply that it can be proper to assert falsehoods and treat them as reasons for action. You might be *excused* for telling others falsehoods or for acting on faulty intelligence, but it isn't *proper* to do so.⁶

¹ Hawthorne and Stanley (2008), Unger (1975), and Williamson (2005).

² Adler (2002), DeRose (2002), Slote (1979), Sutton (2005), Turri (2011), Unger (1975), and Williamson (2000).

³ Fantl and McGrath (2009), Gerken (2011), Gibbons (2010), Littlejohn (2009, 2012), Locke (MS), and Neta (2009).

⁴ Douven (2006), Kvanvig (2009), Lackey (2007), Littlejohn (2012), and Sutton (2005).

⁵ In addition, those who defend knowledge norms (e.g., Williamson (2000) and Hawthorne (2004)) often criticize justification accounts for their failure to handle lottery cases and Moore's Paradox. See Kvanvig (2009), Littlejohn (2010), and McGlynn (Forthcoming) for a response.

⁶ DeRose (2002), Hawthorne and Stanley (2008), and Williamson (2000) all maintain that false assertions are excused, not warranted. Gerken (2011), Kvanvig (2009), and Lackey (2007) all argue that the appeal to excuses will not help fend off the apparent counterexamples. There is a tendency in the epistemology literature to think that the

Or, so say defenders of KPR and KA. For their part, those who defend JPR and JA often say that cases of reasonably held false beliefs show that KA and KPR are too strong. If you told them that it was safe to cross and they fell to their deaths, that might be *regrettable*, but you did nothing *wrong* in telling them that it was safe to cross. Both sides think that cases of reasonably held false beliefs constitute counterexamples to their opponents' views.

Which side is right? Neither. The right response to these cases is not to say that the justification accounts are too weak. The right response is not to say that these cases of reasonable false belief show that knowledge doesn't play a central role in normative appraisal. The right response is to reject those accounts of justification that imply that reasonably held false beliefs constitute counterexamples to KPR and KA.

In this paper, we shall focus on a specific kind of reasonably believed falsehood—false normative propositions about what to do or what has moral value. I shall argue that these cases do *not* threaten KPR. What these cases show is that JPR as its ordinarily understood is indefensible. In the next section, I shall offer two arguments for JPR and distinguish between two approaches to the epistemic norms. I shall argue that reasonably held false beliefs often violate the epistemic norms that govern them and that this shows that a popular approach to epistemic normativity is misguided. The take away point will be that while JPR is correct, KPR is not obviously incorrect.

THE UNIFICATIONIST AND THE RATIONALIST

There is a standard criticism of KPR in the literature.⁷ KPR says that it is improper to treat what you justifiably believe as a reason for action if your belief doesn't constitute knowledge. If, say, you believed that the stuff was gin and your belief was based on good evidence, many epistemologists would say that your belief about the contents of the bottle would be justified even if the stuff in the bottle were petrol. Intuitively, it seems odd to say that your belief about the contents might be justified and then say that it would be improper to treat what you believe as a reason for action.⁸

There are two ways of fleshing out the intuitive worry. First, consider two putative norms of theoretical reason:

paradigm case of an excuse is something like insanity or brain washing. (See, for example, Audi (2001) and Wedgwood (2002).) The thought seems to be that excuses exculpate by showing that the agent's capacities for rational control have been undermined. I think that this is a mistake. Following Strawson (1965) and Gardner (2007), I think it's important to distinguish excuses from exemptions or denials of responsibility. In offering an exemption, you try to remove blame by showing that the agent's capacities for rational control have been undermined. In offering an excuse, you try to remove blame by showing that the agent's rational capacities are intact. We should reject the standard reason for rejecting the appeal to excuses. I defend this view in my (2012). The arguments in this paper lend further support to the idea that the rational or reasonable is *not* the mark of the right or the permissible.

⁷ Brown (2008), Madison (2010), Neta (2009), and Littlejohn (2012) argue that Gettier cases cause trouble for KPR and KA.

⁸ You can find versions of this criticism in Douven (2006), Fantl and McGrath (2009), Gerken (2011), Kvanvig (2009), Lackey (2007), Littlejohn (2012), Locke (MS), Neta (2009), and Smithies (2012).

JTR: Whether it is proper to treat p as a reason for belief depends upon whether you justifiably believe p .

KTR: Whether it is proper to treat p as a reason for belief depends upon whether you know p .

It seems that justification is closed under known entailment. If you justifiably believe, say, that the stuff is gin and know that gin is not petrol, you should be able to justifiably infer that the stuff is not petrol. If, however, it is possible to justifiably believe what you don't know and you don't know that the stuff is gin, KTR seems to imply that it is permissible to believe that the stuff is gin and impermissible to conclude that it's not petrol even though you know that the stuff cannot be petrol if it's not gin. It seems that the right to believe, however, comes with further rights. Among these rights is the right to treat what you believe as a reason for belief. If it didn't come with this right, we'd be at a loss to explain why justification seems to be closed under known entailment. Since it seems you shouldn't say of someone that it was proper for her to believe p but improper to believe p 's obvious consequences on the basis of p , it seems JTR must be the norm of theoretical reason.

If JTR is the epistemic norm that governs theoretical reason, what should we say about the epistemic norms that govern practical reason? Suppose that KPR rather than JPR governed practical reason. Theoretical reasoning aims to settle questions about what is true and comes to its conclusion when you form a belief.⁹ Practical reasoning aims to settle questions about what to do and comes to its conclusion when you form an intention or perform an action. Suppose that you justifiably believed that the stuff was gin, but you were mistaken about that. Suppose that you knew that you ought to serve the guests gin if you can. In the case described, you could properly settle the question whether it is true that you should serve the guests gin by believing that you should and treating the (apparent) fact that the stuff was gin as a reason for so believing. In the case described, you would be obliged to refrain from settling the question whether to serve the guests gin by treating the (apparent) fact that the stuff was gin as a reason for intending or acting. *This* is why we cannot say that practical reasoning is governed by KPR rather than JPR. From the epistemic point of view, whatever reason there is not to treat the (apparent) fact that p as a reason to ϕ is a reason not to treat the (apparent) fact that p as a reason to believe that you ought to ϕ .

The argument for JPR rests on two plausible thoughts. The first is that a common norm has to govern practical and theoretical reasoning. (To say otherwise, you have to say that it would be improper to settle the question as to whether to ϕ even when you have properly settled the question as to whether you should ϕ .) The second is that a justification norm has to be the norm of theoretical reason. (To say otherwise, you have to say that what justifies settling the question whether p by believing p does not justify settling the question whether q even when you know that q follows from p .) If these points are both correct, then the Unity Thesis must be true:

Unity Thesis: You cannot justifiably judge that you should ϕ unless you have justification to ϕ .

Let's say that the *unificationists* accept the Unity Thesis.¹⁰ The unificationists think that the demands of practical and theoretical reason are unified in such a way that

⁹ Hieronymi (2005).

¹⁰ Gibbons (2009, 2010), Littlejohn (2012), and Smithies (2012) defend the Unity Thesis.

your overall obligation is never to act against your own justified judgment. Much in the way that it seems that any reason that counts against acting would count against intending to act, the reasons that count against intending to act would count against judging that you should act.

The second argument for JPR is the *rationalist's* argument. The rationalists think that the mark of the right, the permissible, or the justified is the rational or the reasonable:

Rationality Thesis: You have justification to ϕ if it would be rational for you to ϕ .¹¹

They object to KPR on the grounds that KPR assumes some sort of objectivist account of obligation. An objectivist about obligation will say that your obligations can sometimes depend upon facts that you are non-culpably ignorant of or mistaken about and KPR commits you to such an account of obligation because it says that these sorts of facts can determine whether you are under any obligation to refrain from treating p as a reason for action.

If the Rationality Thesis is correct, the rational belief or the reasonable belief is the justified belief. Were a rationally held belief of yours to be unjustified, you would be obliged to refrain from believing what you did. As it would be reasonable to retain your belief, however, you would be permitted to retain your belief and so your belief would be justified. Suppose that you justifiably believed that you should ϕ . Your obligation could not be to do other than ϕ as doing so would be unreasonable. Thus, the Rationality Thesis entails the Unity Thesis and JPR.

The Rationality Thesis entails the Unity Thesis, but the converse entailment does not hold. The unificationist can, if she wishes, accept an objectivist account of epistemic obligation such as KPR. Combined, KPR and the Unity Thesis entail that it's only proper to believe what you know and that it's impossible to justifiably believe what you don't know.¹² While the considerations that support the Unity Thesis support JPR, they don't necessarily support the rejection of KPR.

I suspect that many epistemologists are rationalists. If this is so, I expect that the rationalist's argument would resonate with them. Recall Cohen's new evil demon objection to reliabilism:

Imagine that unbeknown to us, our cognitive processes (e.g., perception, memory, inference) are not reliable owing to the machinations of the malevolent demon. It follows on a Reliabilist view that the beliefs generated by those processes are never justified. Is this a tenable result? I maintain that it is not.¹³

Anticipating that the reliabilists will say that these deceived subject's beliefs are rationally held without being justified, Cohen says this response to his objection fails because 'justified' and 'rational' are "virtual synonyms". His objection to reliabilism, he says:

¹¹ Gibbons (2009, 2010) and Smithies (2012) defend the Rationality Thesis. Although they don't explicitly endorse it, you'll see places where Audi (2001), Cohen (1984), Fantl and McGrath (2009), Huemer (2007), Kvanvig (2009), and Steup (1999) suggest that they take the rational or the reasonable to be the mark of the permissible.

¹² See Sutton (2005) and Unger (1975) for defenses of this view.

¹³ Cohen (1984: 282).

hinges on viewing justification as a normative notion. Intuitively, if S's belief is appropriate to the available evidence, he is not to be held responsible for circumstances beyond his ken.¹⁴

To think that normative notions have to be understood in this way is just to endorse the rationalist view.

The rationalist rejects KPR because of its objectivism about obligation. It doesn't just reject objectivism about obligation in the epistemic domain, but in all domains. Recall from earlier the gin and tonic case. If you think, as I do, that you shouldn't poison your friends, you would say that the best you could hope for if you poison your friends is an excuse for you and a speedy recovery for your friends. As a unificationist, I think that one of the advantages KPR has over its rivals is that it says that just as you shouldn't poison your friends by serving them petrol, you shouldn't treat the (apparent) fact that the stuff was gin as a reason for serving. It tells you not to act and it tells you not to form the beliefs that rationalize wrongdoing. The rationalist rejects the idea that your obligation in this case is to refrain from poisoning your friend and would reject the suggestion that should accommodate the intuition that it's wrong to poison by adopting an objectivist account of epistemic obligation. In this passage, Fantl and McGrath explain why they think you've acted rightly when you poison your friend in the gin and tonic case:

... It is highly plausible that if two subjects have all the same very strong evidence for *my glass contains gin*, believe that proposition on the basis of this evidence, and then act on the belief in reaching to take a drink, those two subjects are equally justified in their actions and equally justified in treating what they each did as a reason, even if one of them, the unlucky one, has cleverly disguised petrol in his glass rather than gin. Notice that if we asked the unlucky fellow why he did such a thing, he might reply with indignation: 'Well, it was the perfectly rational thing to do; I had every reason to think the glass contained gin; why in the world should I think that someone would be going around putting petrol in cocktail glasses!?' Here the unlucky subject ... is not providing an excuse for his action or treating what he did as a reason; he is *defending it* as the action that made the most sense for him to do ... He is providing a justification, not an excuse.¹⁵

As they see it, the mark of justified or the right action is what it makes sense to do and what it makes sense to do depends upon your evidence, not facts that are obscure to you.

Fantl and McGrath's treatment of the gin and tonic case is in keeping with the rationalist view that the facts that determine an agent's obligation (epistemic or moral) are limited to those facts that determine whether an agent's attitudes or actions are reasonable or rational. If, as I suspect, this rationalist view is largely responsible for the direction that the literature on epistemic norms has taken, then I think a careful

¹⁴ Cohen (1984: 282).

¹⁵ Fantl and McGrath (2009: 125).

examination of this view is called for. I shall argue that the view leads to a muddled understanding of obligation and responsibility.

A RATIONALIST ACCOUNT OF OBLIGATION

Because the rationalists reject objectivist accounts of obligation, they need a sensible subjectivist alternative. They need a view on which the justification of action according to which the reasonable belief about what to do is a proper basis for action. For this reason, they have to deny that the facts that bear on what to do differ from those that determine whether our judgments are rationally held. The most sophisticated subjectivist view of obligation is the prospectivist view.¹⁶ The prospectivist says that an agent's obligation is to do what is prospectively best (i.e., what would maximize expectable value). The expectable value of an option is determined by the probability of its possible outcomes and the probable value associated with these outcomes. An agent's evidence determines the probability that an outcome will eventuate and the probable value associated with the various possible outcomes. For the purposes of this discussion, I shall assume that the agent's evaluative evidence is provided by her intuitions and beliefs. This seems like the natural view for any prospectivist to adopt and it fits nicely with the rationalist idea that the facts that determine overall obligation are limited to those that determine what it would be reasonable for the agent to believe.¹⁷

The prospectivist view is similar to the view that agents ought to maximize expected value.¹⁸ These views agree two agents might be under different obligations if they have different empirical evidence. The reason the rationalist should not characterize an agent's obligations in terms of the notion of expected value is that expected value is a function of the objective value of an outcome. Facts about value might be obscure to an agent. Her evaluative evidence determines what's reasonable for her to think about value. Thus, the account that characterizes an agent's obligations in terms of expected value would deny that there is a tight connection between the agent's rational judgments about what to do and the agent's obligations.

An example should help illustrate these points and explain why the prospectivist view is initially quite attractive. Some hedonists will say that death is bad for us (when it is) because it deprives us of something valuable, a future in which there is a greater

¹⁶ Zimmerman (2008).

¹⁷ This seems like the natural view for the prospectivist. Conee and Feldman (2004) and Wedgwood (2010) defend accounts of evidence in the spirit of this proposal. The prospectivist account says that your overall obligation is to ϕ iff ϕ -ing is the prospectively best option. If we ranked options in terms of the agent's non-evaluative evidence and construed the agent's evaluative evidence as consisting in independent facts about value (e.g., those believed or those known) we wouldn't have the materials to rank the agent's options in such a way that the agent would be under many (or any) obligations. The agent would be permitted to fiddle while Rome burned if the agent wasn't in cognitive contact with the right evaluative facts. Moreover, the agent with a defective sense of what's valuable could follow her (apparent) evidence scrupulously and end up with rational judgments about what to do that didn't correspond to facts about what would be prospectively best. So, the prospectivist needs a subjective approach to evidence in the prospectivist account of overall obligation.

¹⁸ Jackson (1994).

balance of hedons to dolors.¹⁹ Some critics of hedonism say that death is bad for us (when it is) because it deprives us of something that we care about, a future in which there is a greater balance of hedons to dolors.²⁰ These accounts don't necessarily disagree about what makes some experience good for us. They disagree about the kinds of relations we have to bear to these possible future experiences for us to be deprived by losing out on them. If entities that lack the capacity to care for their futures cannot be deprived when painlessly killed, then death might be bad for us even if it is not bad for an embryo. There are cases in which it seems intuitive to say that the destruction of an embryo is not bad for the embryo and cases in which it does seem intuitive to say that allowing the embryo to develop into a creature that has the capacity to care about things is not good for the embryo. The capacity to care condition is not without its intuitive support. Still, there are cases where it seems rather intuitive to say that the capacity to care is not required for a deprivation to be bad. I expect that many readers will have some intuitions/evidence that supports views that incorporate the capacity to care condition and some intuitions/evidence that supports views that don't.

If you had a boat, a flood, and three creatures trapped on the roofs of two barns, you would have to make a decision about what to do. If you steer your boat to the first roof, you can save Doris. Edith and Elsie would perish. If you steer your boat to the second roof, you can save Edith and Elsie. Doris would perish. Doris and Edith are two equally healthy young women. Their futures would be equally good for them. Elsie is a cow. Her future experiences would be uniformly positive, but she lacks the capacity to care about her future. What to do?

If the capacity to care condition is correct, the expected value of heading to the first barn is the same as the expected value of heading to the second. If your evidence supported hedonism to some degree, however, the expected value would differ. Intuitively, it seems that a perfectly good way to reason to a decision about what to do is to take account of the fact that your confidence is rationally divided between these two views. It seems you can properly reason that since it would be just as good to go left as right whichever view of well-being is correct and better to save Edith and Elsie if the hedonist view is correct that you should save Edith and Elsie. A virtue of the prospectivist view is that it seems to be the one view on which good reasoning could take account of these facts. If obligation is determined by expected value or objective value, such reasoning would transparently unsound.

There are two important points to bear in mind. The first is that the prospectivist view vindicates the kind of reasoning sketched above because it allows that the agent's own uncertainty can partially determine the deontic status of the agent's options. The second is that the rationalist needs to adopt something akin to the prospectivist view because they'll want to say that the agent's evaluative evidence plays the dual role of rationalizing judgments about what to do and determining what an agent's obligations are.

TWO APPROACHES TO EXCULPATION

Mistakes and ignorance can exculpate. If you reasonably thought that the petrol in the bottle was gin, you could not be blamed for serving the stuff to your friend. On objectivist accounts of obligation, it might be wrong to serve your friend the stuff, but

¹⁹ See Bradley (2009).

²⁰ Velleman (2000).

you can be excused for doing so if you were non-culpably ignorant of the fact that the stuff was poisonous. The objectivist might treat all cases of non-culpable ignorance as being on par, but the objectivist is not committed to that. Instead, they can (and should) say that when it comes to moral responsibility there is a difference between cases of evaluative and non-evaluative ignorance:

The Asymmetry Thesis: Non-culpable factual mistake and ignorance will typically excuse the agent's behavior, but non-culpable evaluative mistake and ignorance will typically not.

On an attractive approach to moral responsibility, you're responsible for wrongdoing when your actions manifest a certain kind of moral unresponsiveness. Specifically, you can be blamed for *de re* moral unresponsiveness, a failure to show proper sensitivity to the concerns of morality.²¹ This account, the quality of will account, explains the Asymmetry Thesis. The fact that you poisoned your friend might initially seem to be evidence that you do not respond properly to the prospect of harming your friend. When, however, we learn that you are non-culpably ignorant of the fact that the stuff you served was poisoned, we can see that your actions do not manifest *de re* unresponsiveness. Your actions would display *de re* unresponsiveness if, say, morality cared about whether your friend suffered harm and your actions manifested an indifference or willingness to impose that harm. Excuses do their work by undermining an inference from observations having to do with the objective nature of your deed or the consequences of your action to judgments about the quality of your will.²²

Rationalists agree that factual mistakes and ignorance can exculpate, but they think that the objectivist model is wrong because it implies that facts that you're non-culpably ignorant of can determine what you're obligation is. On the rationalist view, if you act on reasonable but mistaken beliefs about the facts, you wouldn't need an excuse if, say, you brought about bad consequences. If you reasonably believed that the stuff in the bottle was gin, you did nothing wrong by serving your friend a Bernie (i.e., a petrol and tonic with lime on the rocks). Ignorance and mistake exculpate by subverting obligation, not excusing wrongdoing.

The trouble with the rationalist view is that it is incompatible with the Asymmetry Thesis. To see this, we need to shift our focus to cases of normative ignorance and mistake. The quality of will account says that factual ignorance and mistaken factual belief can exculpate by showing that the agent's actions do not manifest *de re* moral unresponsiveness. The account does not (typically) dish out excuses when agents act on mistaken normative beliefs or act from normative ignorance. If the agent acts in full awareness of the facts and engages in wrongdoing, the agent's behavior manifests a willingness to act against morality's concerns. This, the quality of will account says, is precisely what we should blame agents for when they should be blamed for wrongdoing. We blame them for the failure to show proper sensitivity to morality's concerns.

What should the rationalist say about normative ignorance and normative mistake? From the epistemic point of view, mistakes are mistakes. The subject matter of a mistaken belief is not the sort of thing that determines whether a mistake is reasonable. Fit between the belief and the evidence is what determines whether a belief is reasonably held. Because the rationalist is committed to the view that moral beliefs are reasonably

²¹ For an articulation and defense of the view, see Arpaly (2003).

²² Strawson (1962).

held if they fit the evidence (i.e., with the subject's beliefs and intuitions), the rationalist cannot classify the actions rationalized by reasonable judgments about what to do as wrongs. Because of this, they cannot classify them as culpable wrongs.

The rationalists cannot say that we are culpable for engaging in wrongdoing if we are non-culpably ignorant of the facts or their moral significance. As such, they have to reject the quality of will account and its suggestion that we're to blame for *de re* moral unresponsiveness. In denying that we're to blame for failing to show proper sensitivity to morality's concerns, the rationalists needn't deny that we're to blame for failing to show sensitivity to morality. The rationalist might say that we're to blame for failing to show due deference to morality as a source of authoritative or overriding reasons when, say, we decide to pursue our own interests or the interests of those near and dear to us rather than meet our moral obligations. In other words, they might say that we're to blame if our actions display *de dicto* moral unresponsiveness, a willingness to do what we take ourselves to be duty bound not to do. On this account, we might also be responsible for doing what we should have thought we'd be obliged to do. The important point is that this view implies that non-culpable factual and normative ignorance and mistake exculpate.

The rationalist's view doesn't just excuse the inexcusable; it justifies it. Don is your typical television father from the late 50s or early 60s. He loves his children and does what he can to try to keep them happy. He puts money away for his son to go to school. He puts money away for a sailboat for his daughter. Although his daughter has said repeatedly that she wants to go to school, he sees no reason to help her because he sees no reason for women to go to college. Intuitively, he should have put money away for his daughter's education. He should have done so even though there was nothing in Don's evidence that would rationally support his belief that he ought to support his daughter's ambitions in the same way he supports his son's ambitions.

Not only did Don fail in his responsibilities as a father, he is responsible for this failing. He knew that both of his children wanted to go to college and he knew that he could provide for both of them. He was fully aware of the reasons that there were to save for his daughter's education. True, he didn't see these reasons as reasons, but this is why we think that he is sexist, not why we should excuse him.

Because Don shouldn't act like a sexist and his sexism isn't an excuse for his actions, he's culpable for his wrongful behavior. The rationalist doesn't have the resources to deliver the right verdict about Don's responsibilities or Don's responsibility for failing to meet them. Plausibly, the reason that the otherwise virtuous Don didn't do what we now think we should for our children is that Don lacked the evidence he needed to rationally settle the question as to whether to save for his daughter's education in the way that he should have. If the rationalists want to deny this, they have to say that Don's judgments about what to do are somehow less than fully rational. The rationality of these judgments, however, depends upon how he responded to the evidence he had, which we take to be the intuitions and beliefs that he had. There is no good route of reasoning that would take him from those beliefs and intuitions to the judgment that he should treat his children the way that we think he should.

The case shows that the rational or reasonable is *not* the mark of the right, the justified, or the permissible. If it were, Don wouldn't be culpable for wrongdoing as there would be no wrongful act that he might then be blamed for. What would be rational for the morally conscientious agent to believe is determined by the agent's evaluative evidence and empirical evidence. If the rationalist view is correct, the same facts

determine what's reasonable to do and what's permissible to do. We know that the rationalist view isn't correct because we know that the facts that determine what's rational for Don to believe don't determine what Don is permitted to do. Even if we were to say that there is no real gap between the rational and the justified in the case of *factual* ignorance or mistake, there is a yawning gap between the two in cases of *normative* ignorance and mistake.

If the objection is sound, what does it show? For one, it shows that the rationalist's view is unsound. It shows that you cannot argue from rationalist assumptions against KPR or appeal to rationalist assumptions in offering an account of justified belief. For another, it shows that if the unificationist argument is sound, JPR requires us to reject all orthodox accounts of justification. All the orthodox accounts say that justification is non-factive. JPR says that justified beliefs are the proper basis for action. If you justifiably but falsely believed, say, that putting money aside for a sailboat was the thing to do, it wouldn't be something you should do and it would be something you'd have the right to do. So, arguing from unificationist assumptions, we should be able to show that the orthodox accounts of justification are wrong. Your normative beliefs might fit the evidence, they might be formed reliably, they might cohere, etc., but they wouldn't be justified if they didn't fit the normative facts.²³

A RATIONALIST RESPONSE

The rationalist has to reject the Asymmetry Thesis. The rationality of normative beliefs don't depend upon whether they (or the beliefs they're based upon) accurately represent the values that determine what the agent's obligations are. Since the rationalist takes the rational judgment about what to do to be a justified basis for action, the most the rationalist can say about the actions the agent performs on the basis of a rational judgment is that it is regrettable. The rationalist cannot say that there's any wrong to excuse, and so she must reject the Asymmetry Thesis and the quality of will account. Whatever it is that renders us blameworthy for the things that we do, it's not mere *de re* moral unresponsiveness. If it's not *wrong* to display such insensitivity, it's not *culpably* wrong to do so. Since this seems to be intuitively wrong, so much the worse for the rationalist view.

This, in brief, is the anti-rationalist argument. What can the rationalists say in response? I think the only thing that the rationalist can do is reject the Asymmetry Thesis. In this section, I shall argue that if the rationalist does this, the rationalist faces some serious problems. By rejecting the Asymmetry Thesis, the rationalist view leads to a skeptical view about moral responsibility.

Suppose you serve your friend a Bernie and the drink nearly kills her. Plausibly, you cannot be blamed for the poisoning if you were non-culpably ignorant of the fact that the stuff you served was poisonous. Everyone should agree to that. The rationalist might then say that this reflects an important point about culpability and ignorance:

²³ The argument isn't intended to show that *all* beliefs must be true to be justified, only that an important class of beliefs cannot be justified if false. Since orthodox accounts of justification say otherwise, Don's belief about what to do for his daughter is a perfectly good counterexample to all the familiar accounts of justification that allow for false, justified beliefs (e.g., reliabilism, evidentialism, phenomenal conservatism, proper-functionalism, etc.).

The Culpability Thesis: If you are ever culpable for acting in ignorance, you have to be culpable for the ignorance in which you act.²⁴

The rationalist can appeal to the Culpability Thesis to try to explain why we should reject the Asymmetry Thesis. Just as you cannot be culpable for poisoning your friend if you were non-culpably ignorant of the fact that the stuff you served was poisonous, Don cannot be culpable for his wrongdoing if he's non-culpably ignorant of the reasons for treating his daughter differently. By hypothesis, Don has responded to his evidence responsibly and rationally, so he cannot be culpable for his failure to believe that he ought to treat his daughter differently.

The rationalist's only line of response is to appeal to the Culpability Thesis to try to undercut the Asymmetry Thesis. Unfortunately, the Culpability Thesis leads rather quickly to widespread skepticism about responsibility. This is because the Culpability Thesis entails the Origination Thesis:

The Origination Thesis: Culpable wrongdoing can be traced back to a case of clear-eyed akrasia (i.e., a case in which the agent acts in the belief that the relevant action is wrong).

According to the Culpability Thesis, Don can't be culpable for his wrongful behavior if he doesn't believe it to be wrong unless he's culpable for his ignorance. What would it take for him to be culpable for his ignorance? The rationalist cannot say he's culpable for his ignorance *simply* because his ignorance manifests, say, his sexism. To be culpable for his ignorance, he has to be culpable for the way he's managed his opinion. And, as Rosen notes, to be culpable for the failure to believe, Don must be culpable for failing to take the appropriate steps in forming his beliefs. To be culpable for that, however, there must be something that he's done or that he's left undone that he's culpable for that's responsible for his ignorance (i.e., his failure to form a true belief about the relevant normative matter). What could that be? It could not be further ignorant behavior on his part. If it were, the argument would apply again. Thus, once the rationalist endorses the Culpability Thesis, the rationalist has to endorse the Origination Thesis.

When combined with the plausible premise that we rarely act in the belief that our obligation is to do otherwise, the Origination Thesis leads to an implausible form of skepticism about moral responsibility. Since the rationalist is committed to the Culpability Thesis, they can try to show that these theses are independently plausible and that the consequences aren't quite so bad as they might first appear. In this section, I will argue for two points. The first is that the arguments offered in support of the Culpability Thesis aren't persuasive. The second is that the rationalist view actually undermines the quite modest thesis that you must endorse if you think anyone is ever responsible for anything:

The Trivial Culpability Thesis: If an agent knowingly engages in wrongdoing and is not under duress, the agent is culpable for his wrongful behavior.

Rosen recognizes that the implications of the Culpability Thesis might seem *prima facie* implausible, but thinks that it is independently plausible:

²⁴ It also seems that the rationalist is committed to the Culpability Thesis. If you reject the Asymmetry Thesis and believe that non-culpable factual ignorance exculpates, you have to believe that non-culpable evaluative ignorance exculpates. If you didn't, you would reject the Asymmetry Thesis.

When we engage the examples [of moral ignorance and mistake] in imagination, bearing it fully in mind that the agent is not responsible for his moral ignorance, then our capacity to blame is neutralized by this very thought. Moreover, this is not simply a psychological observation. When we find ourselves unwilling to blame the agent who acts from blameless ignorance, it is because we have come to think that it would be a mistake to blame him.²⁵

If you press Rosen to explain why non-culpable ignorance exculpates, he says that it would be unfair to blame someone for his action if the agent lacks the general capacity to appreciate and to act on moral reasons there are to do otherwise. Extending this point, he says:

Just as it is unfair to blame someone for doing what he is in fact entitled to do, it is unfair to blame someone for doing what, through no fault of his own, he *takes himself* to be entitled to do. Take [Don] at the point at which he has formed the view that it is perfectly all right for him to deny his daughters certain opportunities. So far, by hypothesis, he is blameless. Now he does it. How can you blame him? How can you expect him to do otherwise given what he blamelessly believes?²⁶

Rosen is hardly alone in thinking this. Even some of his critics agree that the crucial question to ask about Don is whether he could be reasonably (and hence fairly) expected to do something to remedy his ignorance.²⁷ If we assume that Don displays a sufficient degree of concern about meeting morality's demands (whatever they are) and scrupulously fits his beliefs to the evidence, there are no steps they could have reasonably taken to remedy their ignorance. Rosen thinks that this should be exculpatory.

Zimmerman offers a different defense of the Culpability Thesis, one that has to do with control. We've seen already that you cannot be held responsible for acting out of *factual* ignorance unless you are culpable for the ignorance in which you act, but what about *moral* ignorance? Zimmerman thinks that you cannot be culpable for something unless you're in control of that thing and that you cannot be *directly* culpable for something unless you're *directly* in control of that thing. We aren't directly in control over whether we're ignorant. If we're ever to remedy our ignorance, we have to do so by means of something else we have control over (e.g., considering a thought experiment, asking a guru, running through an argument, etc.). If this is right and we're never directly in control over whether we act from ignorance, we're never directly responsible for acting from ignorance whether the ignorance has to do with matters of morality or matters of fact.²⁸

²⁵ Rosen (2003: 71).

²⁶ Rosen (2003: 74).

²⁷ See FitzPatrick (2008) and Levy (2011). Levy is right that FitzPatrick is too optimistic in thinking that our baddies could have taken steps to discover the moral truth. Both are mistaken in thinking that it matters whether they were.

²⁸ Zimmerman (2002, 2008) offers this defense of the Culpability Thesis.

Neither defense is compelling. Rosen's argument shouldn't persuade anyone who doesn't already reject the quality of will account. Whether you're culpable for your *ignorance* (i.e., for your lack of belief) depends upon whether your ignorance shows that you are properly concerned with the values that matter to epistemic assessment (i.e., the pursuit of truth and avoidance of error). In keeping with the quality of will account, you are not culpable for your error if you are properly sensitive to that which has epistemic value (i.e., if you follow your evidence). However, whether you are culpable for your *action* depends upon whether your actions show that you have proper concern for the values that matter to moral assessment. The person of reasonable prudence can form beliefs that show that she is properly responsive to epistemic reasons even if her actions manifest *de re* moral unresponsiveness. If that's so, there's no reason to think that Don cannot be culpable for his sexist behavior simply because he's not culpable for failing to form the belief that he should have done otherwise.

Zimmerman's argument for the Culpability Thesis assumes that what you're directly responsible for is limited to what you have direct control over. This is also a mistake. To say that you are *directly* responsible for something is to say that you are responsible for it and that your responsibility does not derive from your responsibility for something else. Those who accept the quality of will account will say that the reason that you are culpable for wrongdoing when you are is that you are not properly sensitive to the interests morality wants us to protect. The question as to whether you can directly remedy the defects in your will is neither here nor there.

It doesn't seem to me that those who accept the Culpability Thesis have offered an adequate defense of it. If the Culpability Thesis is unmotivated, the rationalist hasn't given us any good reason to reject the Asymmetry Thesis. Someone might succeed where others have failed, so let's suppose that the Culpability Thesis is correct and see whether we can live with its consequences. Some of us would say that the skeptical consequences of the Origination Thesis are bad enough and reject the rationalist view on the grounds that it implies that the locus of original responsibility cannot be anything but a case of clear-eyed akrasia. As bad as this is, I think the rationalist's predicament is actually much worse than this. I don't think the rationalist can make sense of how somebody could be culpable even in these cases.

If the rationalist is committed to the Culpability Thesis, she needs an account of culpability that accommodates it. It seems that two options are available to her. First, she might say is that we're culpable for wrongdoing when we act in the *belief* that we're doing wrong. The action is culpable because it shows that the agent doesn't show due respect for morality as a source of reasons. Second, she might say that we're culpable for wrongdoing when we act in the *correct belief* that we're doing wrong. The action is culpable because it shows that the agent doesn't show due deference to morality and shows that the agent's concerns aren't morality's concerns. Neither approach is at all plausible.

Let's consider the first approach first. It has a number of objectionable features. First, the thesis that we can be blamed for failing to show due deference to morality is subject to counterexamples involving inverse akrasia.²⁹ When Huck Finn helped Jim escape, he did so in the belief that what he did was wrong. He thought that it was morally wrong to help runaway slaves. If we are culpable for the failure to show due deference, we would have to say that Huck was culpable. To show due deference to

²⁹ See Arpaly (2002).

morality, you have to respond to morality's apparent demands by treating them as authoritative. Acting against your conscience would be a failure to show due deference and while Huck did that, he was praiseworthy for it.

Second, since the view assumes that you are not directly responsible for the failure to show proper sensitivity to morality's concerns, we have to ask why we would be responsible for the failure to show due deference to morality if there is no specific concern of morality that we can be held responsible for being insensitive to. The answer must be that morality has a kind of rational authority by virtue of which failures of due deference are inculcating in a way that the failure to show sensitivity to moral reasons wouldn't be. This generates an explanatory worry. I cannot see any reason to think that there would be an explanation of the rational authority of morality that would explain why failure to show due deference is inculcating if insensitivity is never itself inculcating on its own. If morality's reasons don't have rational authority, why would morality?

Third, it is not clear that the rationalist is entitled to the assumption that morality has the rational authority it must for a failure of due deference to be a locus of original responsibility. Consider the legal case. The law aims to protect a number of interests. An agent that threatens those interests might do so because she is ignorant or mistaken about the facts. In such cases, it seems inappropriate to punish because the agent has not shown herself to be willing to injure any of the interests that the law aims to protect. On this point, the rationalists and the rest of us can agree. What about cases of impossible attempts in which the agent tries to perform an action she believes mistakenly to be illegal? Does such a mistake of law inculcate? No, it seems not. As with the case of mistake of fact where an agent unwittingly threatens a legally protected interest without showing that she is willing to harm this interest, the agent who attempts to do what she falsely believes to be illegal has not shown herself to be willing to injure any of the interests that the law protects. This assumes, of course, that the law's interests do not include an interest in having citizens show due deference to the law. The law can properly protect us from those who would deprive us of property or exploit children, but not protect us from those who are willing to break a law in circumstances where no further legally protected interest is at stake. If it did otherwise and its list of protected interests included deference to the institution of law itself, I think we would have little respect for the institution. It would be willing to punish agents who were unwilling to threaten any of its other interests simply when such agents followed all of its regulations for the simple failure to respect authority as such. That seems outrageous.

An institution that would punish or sanction someone who fails to respect its authority without having failed to show proper regard for the interests that the authority is dedicated to protecting is defective. Its sanctions and punishments would be the empty exercise of authority and it would not deserve the kind of respect that moral rationalists have claimed morality deserves. It would be an empty exercise of authority for morality to sanction you for the failure to show due deference when morality acknowledges that you showed proper sensitivity. Just as a just system of laws will not regard mistake of law as inculcating, a just moral framework will not take failure to show due deference to be the locus of original responsibility.

At this point, the rationalist might agree that acting against your own moral judgments is not inculcating unless your judgment is correct. This would allow the rationalist to avoid the three problems just discussed, but it doesn't save the rationalist view. We can see why if we think about a concrete case. The rationalist view says that rational judgments about what to do can justify acting in line with those judgments. The

rationality of these judgments depends upon an agent's evaluative evidence (i.e., her intuitions and beliefs). With this picture before us, it seems to me that the following is a plausible conjecture: with a handful of exceptions there is for each reading assigned in your standard course on contemporary moral problems a possible configuration of mental states such that a subject could rationally believe the conclusion of that piece for roughly the reasons set out in that piece. The rationality of the moral judgments will depend upon whether there is some sound deliberative route that takes you from the evaluative evidence to the relevant conclusion about responsibilities. (Remember that we typically can only undo the work that should be undone by changing the evidence, not simply by pointing out that the reasoning that led to the conclusion was faulty. We do not typically take these authors to be less than fully reasonable for failing to dream up the intuition pumps that later authors devise to frustrate their efforts.)

With this conjecture in mind, consider Aquinas' claim that essentially non-procreative sex is morally defective. He might have thought that if such sex were required for the preservation of the species or for shutting down the Doomsday device, it would be acceptable, but his arguments seem to require him to believe that it would be wrong for two men to have sex with each other because they had finished a bottle of wine and their cable was on the fritz. We might imagine that Adam and Stephen shared Aquinas' intuitions and relevant beliefs, so there was no sound deliberative route that would take them to the conclusion that it would be acceptable for them to do what they were about to do. The judgment that we really shouldn't do this has never kept us in our clothes. You know how this ends.

What should we say about Adam and Stephen? What we should say is that their actions were not for the prospectively best. Given their evaluative evidence, abstaining ranks higher than indulging. The rationalists have to say that if any action is ever impermissible, their lusty romp was impermissible. We should not, however, say that their actions were impermissible. The view delivers the wrong deontic verdict about the case, but that's not my concern here. My concern has to do with culpability? Should we say that Adam and Stephen are culpable for acting when they acted in the belief that what they were doing was wrong?

Remember that their belief is a reasonable one. It is the belief that they formed by following the evidence they had. Adam and Stephen believe that they have done wrong on the basis of evidence. This evidence, the rationalist says, justifies their beliefs *and* ensures that their beliefs are correct. So, the rationalist doesn't just think that this is another case of clear-eyed akrasia, but also a case in which the agents *knowingly* do wrong. If they knowingly do wrong, the Trivial Culpability Thesis says that this is a case where Adam and Stephen are culpable. The rationalist faces a dilemma. If the rationalist thinks that we shouldn't blame Adam and Stephen, they must reject the Trivial Culpability Thesis. If the rationalist endorses the Trivial Culpability Thesis, then they have to blame Adam and Stephen and say that their ignorance is inculpating.

The problem with the first horn of the dilemma is obvious. If you can't blame someone for knowingly doing wrong when he's not under duress, it doesn't seem you can blame anyone for anything. So, rejecting the Trivial Culpability Thesis pushes us towards total skepticism concerning moral responsibility. If, on the other hand, the rationalist thinks we should blame Adam and Stephen, then the rationalist has to think that there's some sense in which Adam and Stephen have acted badly. I suppose that most rationalists would grasp this second horn of the dilemma.

If the rationalist grasps the second horn of the dilemma, they have to say that Adam and Stephen have acted badly. In what sense have they done that? The rationalist cannot say that they acted badly simply because they failed to show due deference to morality. This would be to revert back to the view of responsibility we've just rejected. If we blame Adam and Stephen for the failure to show due deference in a case where we know they haven't failed to show proper sensitivity to any of morality's concerns, it would be fetishistic of us to desire that they not act badly in the relevant sense. We would express the desire that agents show due deference to morality on those very occasions where we judge that the agents were sensitive to morality's concerns. We should no better than to blame those we know have been sensitive to morality's concerns. Thus, it seems that the rationalist would have to say that Adam and Stephen have acted badly (in part) because they've shown *de re* moral unresponsiveness. But, by our lights, Adam and Stephen haven't failed to show proper sensitivity to morality's interests.

Because the rationalist has to reject even the Trivial Culpability Thesis, the rationalist cannot avoid wholesale skepticism about moral responsibility. When I said earlier that that the rationalist view offers a muddled account of moral responsibility, this is the muddle I had in mind.

CONCLUSION

Let me briefly recap. The rationalist rejects KPR on the grounds that it is committed to an objectivist account of obligation. If, as the rationalist claims, it's proper to treat reasonable beliefs as reasons for action, this must be true of reasonably held *normative* beliefs. Because the rationalist thinks that any reasonably held normative judgment is a justified basis for action, the rationalist view implies that normative ignorance subverts obligation and that it exculpates. Moreover, they're faced with the unpalatable choice between a view on which normative ignorance inculpates and wholesale moral skepticism.

The rationalist view was that there were epistemic constraints on practical normativity. The rationalist can defend the Unity Thesis on the grounds that these constraints would rule out any case in which an agent's obligation is to act against her own proper judgment about what to do. This gets things backwards. There are practical constraints on epistemic normativity and this is why the requirements of practical and theoretical reason are unified. As Hume never said, the standards of theoretical reason are, and ought to be, the standards that determine whether you meet independently determined standards of practical reason. They cannot pretend to any other office.

What does this all tell us about justification and epistemic norms? Recall from earlier the unificationist argument for JPR. It starts from the idea that the justified belief is, *inter alia*, the justified basis for action. To justifiably believe *p*, you must have the right to treat *p* as a reason when trying to settle questions about what's true and what to do. If that's so, the anti-rationalist argument shows that the justification of at least some beliefs (e.g., evaluative beliefs) depend upon whether they fit the facts, not just upon whether they fit the evidence. If the unificationist argument for JPR sketched above was sound, the upshot is *not* that we can use some orthodox conception of justification to tell us what it takes to conform to the norms governing belief. The upshot is that there's little that distinguishes justification from knowledge.

REFERENCES

Adler, J. 2002. *Belief's Own Ethics*. MIT Press.

- Arpaly, N. 2002. *Unprincipled Virtue*. Oxford University Press.
- Audi, R. 2001. An Internalist Theory of Normative Grounds. *Philosophical Topics* 29: 31–45.
- Bradley, B. 2009. *Well-Being and Death*. Oxford University Press.
- Brown, J. 2008. Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning. *Nous* 42: 167-89.
- Cohen, S. 1984. Justification and Truth. *Philosophical Studies* 46: 279–95.
- Coffman, E. J. 2011. Two Claims about Epistemic Propriety. *Synthese* 181: 471-88.
- Conee, E. and R. Feldman. 2004. *Evidentialism*. Oxford University Press.
- DeRose, K. 2002. Assertion, Knowledge, and Context. *Philosophical Review* 111: 167-203.
- Douven, I. 2006. Assertion, Knowledge, and Rational Credibility. *Philosophical Review* 115: 449–85.
- Fantl, Jeremy and Matt McGrath. 2009. *Knowledge in an Uncertain World*. Oxford University Press.
- FitzPatrick, W. 2008. Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge. *Ethics* 118: 589-613.
- Foley, R. 2001. The Foundational Role of Epistemology in a General Theory of Rationality. In A. Fairweather and L. Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Oxford University Press, pp. 214–31.
- Gardner, J. 2007. *Offences and Defenses*. Oxford University Press.
- Gerken, M. 2011. Warrant and Action. *Synthese* 178: 529-47.
- Gibbons, J. 2009. You Gotta Do What You Gotta Do. *Nous* 43: 157–77.
- _____. 2010. Things That Make Things Reasonable. *Philosophy and Phenomenological Research* 81: 335–61.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Hieronymi, Pamela. 2005. The Wrong Kind of Reasons. *Journal of Philosophy* 102: 37–57.
- Ichikawa, J. 2012. Knowledge Norms and Acting Well. *Thought* 1: 49-55.
- Jackson, Frank. 1991. Decision-Theoretic Consequentialism and the Nearest and Dearest Objection. *Ethics* 10: 461–82.
- Kvanvig, J. 2009. Assertion, Knowledge, and Lotteries. In P. Greenough and D. Pritchard (ed.), *Williamson on Knowledge*. Oxford University Press, pp. 140–61.
- Lackey, Jennifer. 2007. Norms of Assertion. *Nous* 41: 594–626.
- Levy, N. 2011. *Hard Luck*. Oxford University Press.
- Littlejohn, C. 2009. On Treating Something as a Reason for Action. *Journal of Ethics and Social Philosophy*.
- _____. 2010. Moore’s Paradox and Epistemic Norms. *Australasian Journal of Philosophy* 88: 79-100.
- _____. *Justification and the Truth-Connection*. Cambridge University Press.
- Locke, D. Manuscript. Knowledge-Free Decision Theory.
- Madison, B. 2010. Is Justification Knowledge? *Journal of Philosophical Research* 35: 173-92.
- McGlynn, A. Forthcoming. Believing Things Unknown. *Nous*.
- Neta, R. 2009. Treating Something as a Reason for Action. *Nous* 43: 684-99.
- Piller, C. 2007. Ewing’s Problem. *European Journal of Analytic Philosophy* 3: 43-66.
- Rosen, G. 2004. Skepticism about Moral Responsibility. *Philosophical Perspectives* 18: 295-313.

- Sher, G. 2007. *In Praise of Blame*. Oxford University Press.
- Smith, M. 1994. *The Moral Problem*. Blackwell.
- Smithies, D. 2012. The Normative Role of Knowledge. *Nous* 46: 265-88.
- Steup, M. 1999. A Defense of Internalism. In L. Pojman (ed.), *The Theory of Knowledge: Classical and Contemporary Readings*, 2nd edition. Wadsworth Publishing, pp. 373–84. 35:
- Strawson, P. F. 1962. Freedom and Resentment. *Proceedings of the British Academy* 48: 1–25.
- Sutton, J. 2005. Stick To What You Know. *Nous* 39: 359-96.
- Turri, J. 2011. The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy* 89: 37-45.
- Unger, P. 1975. *Ignorance*. Oxford University Press.
- Velleman, D. 2000. *The Possibility of Practical Reason*. Oxford University Press.
- Wedgwood, R. 2002. Internalism Explained. *Philosophy and Phenomenological Research* 65: 349–69.
- _____. 2010. The Moral Evil Demons. In R. Feldman and T. Warfield (ed.), *Disagreement*. Oxford University Press, pp. 216-47.
- Weiner, M. 2005. Must We Know What We Say? *Philosophical Review* 114: 227-51.
- Whiting, D. Forthcoming. Stick to the Facts: On the Norms of Assertion. *Erkenntnis*.
- Williamson, T. 2000. *Knowledge and its Limits*.
- _____. 2005. Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge. *The Philosophical Quarterly* 55: 213: 35.
- Zimmerman, M. 2002. Controlling Ignorance: A Bitter Truth. *Journal of Social Philosophy* 33: 483-90.
- _____. 2008. *Living with Uncertainty*. Cambridge University Press.