King's Research Portal

DOI:

[10.1093/oso/9780198779667.003.0003](https://doi.org/10.1093/oso/9780198779667.003.0003)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Alvarez, M., & Littlejohn, C. (2017). When Ignorance is No Excuse. In J. Wieland, & P. Robichaud (Eds.), *Responsibility: The Epistemic Condition* Oxford University Press.

<https://doi.org/10.1093/oso/9780198779667.003.0003>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

When Ignorance is No Excuse

Maria Alvarez and Clayton Littlejohn

1. Introduction

Ignorance and mistaken belief are often perfectly good excuses. If you didn't know that arsenic had been added to the sugar you should be excused for putting what you did in the tea. If you mistakenly believed that the stuff in the petrol can was petrol, you should be excused for filling your friend's car with Bernard's missing gin. There are interesting debates about whether non-culpable factual ignorance and mistake subvert obligation, but little disagreement about whether non-culpable factual ignorance and mistake exculpate.¹

When people first think about these kinds of cases, they are sometimes tempted to say two things. The first is that these cases show that you are excused for failing to meet your obligations if you fail to meet them because you were non-culpably ignorant of what they were. While not all ignorance or mistake will be exculpating, non-culpable ignorance and mistake should get you off the hook.² The second is that the epistemologists are the ones who will tell us whether we are culpable for our ignorance or mistaken beliefs:³ the standards that determine whether we are culpable for our ignorance and mistaken belief are distinctively epistemic standards. We would not say, for example, that someone is culpable for her beliefs or her ignorance simply on the grounds that the agent's epistemic imperfections explained why she acted badly or failed to meet her obligations. This gets things back to front. If

¹ In many of the interesting cases, the subject isn't ignorant because she is mistaken about the issue in question but because she doesn't have an opinion about it. There might be interesting differences between the ways that these two kinds of epistemic imperfections exculpate. For helpful introductions to this issue, see Arpaly (2001) and Gardner (2007).

² See Gibbons (2013), Rosen (2003, 2004), Wieland (forthcoming), and Zimmerman (2008), for example.

we wanted to explain why you couldn't be blamed for mistakenly believing that you were spooning sugar into the tea, we might say that someone with just your perspective could have reasoned just as you did and formed the same belief after having discharged all of her procedural epistemic obligations. If we wanted to explain why you couldn't be blamed for failing to see that you were filling the tank with gin, we might say that you couldn't be blamed if you couldn't have reached a different conclusion given your perspective on things while meeting your procedural epistemic obligations. If there is no process of good reasoning that takes you from your present beliefs, apparent memories, perceptual experiences, and intuitions to the desired conclusions, it wouldn't be fair to blame you for failing to reach these conclusions. In short, these cases show that a morally culpable failure to meet an obligation requires some sort of culpable epistemic failure.

We think that this conclusion is too sweeping. Our initial cases involved a mistake about or ignorance of some non-moral fact. On the assumption that our agents would not want to poison their guests or damage their friend's car, these agents should probably be excused. Perhaps they should be excused, in part, because their beliefs are just the beliefs we would expect someone to have if they met their procedural epistemic obligations.⁴ What about agents who have all the relevant facts in view but fail to meet their obligations because they do not have the right moral beliefs? If their ignorance of their obligations derives from mistaken moral beliefs or from ignorance of the moral significance of the facts they have in view, should they be excused for failing to meet their moral obligations? It is not obvious that they should. We think that not all epistemic imperfections exculpate equally. In this paper we shall argue that the best non-skeptical accounts of moral responsibility incorporate the Asymmetry Thesis, that is, the thesis that epistemically non-culpable factual ignorance and mistake will excuse in a way that moral ignorance and mistake will not. Ignorance and mistake are

³ For clear statements of this, see Foley (2001), Gibbons (2013), and Rosen (2004).

⁴ We can think of procedural epistemic obligations as having to do with the way that one reasons and revises one's beliefs. Some epistemologists have argued that even if someone argues carefully and in accordance with the canons of good reasoning it is possible to reason from justified premises to unjustified conclusions. For a discussion of the place of procedural epistemic obligations, see Littlejohn (2012).

often perfectly good excuses, but sometimes they do little to nothing to excuse our bad behavior.

2. Setting the Stage

Our stated aim is to defend the asymmetry thesis, the thesis that says that the exculpatory power of our epistemic imperfections depends, in part, upon whether our ignorance or mistaken belief is moral or non-moral.⁵ Our opponents think that a subject might form a reasonable or rational false belief about her obligations on the basis of moral ignorance or mistaken belief. If the agent were to act on this belief or from this ignorance and act impermissibly it would be inappropriate to blame her for her actions even if she failed to see what her obligations were only because of a mistaken moral belief or because of moral ignorance. If Ayn treats her workers as mere means to her ends, say, she might be aware of the demands she makes of them and of the effects that these demands have on their well-being without being aware that it is wrong to demand so much from her workers. It is an interesting question whether this failure to see the moral significance of the facts in view should excuse her when she exploits these poor souls. Depending upon the details, our opponents think that her moral ignorance might just get her off the hook.

We think the opposition is missing something important. In acting from moral ignorance, the subject's actions can manifest *de re* unresponsiveness. We think that it is entirely appropriate to hold people accountable for actions that manifest this kind of unresponsiveness.⁶ We do not see why it is fairer or more plausible to hold people accountable for performing actions that manifest *de dicto* unresponsiveness, so that a man might be culpable for extorting money from vulnerable people if he regards that as immoral but not if he doesn't. It seems strange to us that someone who cares about what morality values (under that description) but does not act with the appropriate

⁵ We will bracket the interesting question as to whether further kinds of normative mistake or ignorance exculpate.

⁶ We side with Arpaly (2001) and Harman (2011). These authors conceive of *de re* responsiveness in terms of desires and beliefs, respectively. We leave that question open.

concern for the things that morality values should escape moral blame. It would be as if morality was more concerned with our concern for its authority - for what carries the label 'immoral', than our concern for its concerns - for what *is* immoral. Since it would be perverse for morality to have such concerns, we feel confident that such concerns couldn't be morality's concerns.

We understand that there are reasonable grounds for doubting the basic picture we wish to defend in this paper, but it seems that the most troubling criticisms of the present approach are criticisms that support a kind of general skepticism about moral responsibility that our opponents in this debate don't accept. Our opponents agree that we can properly hold people accountable for their misdeeds, but they think that moral ignorance and mistake can excuse us when our actions manifest *de re* unresponsiveness so long as they are epistemically blameless. By contrast, we think that the best approach is one that sees *de re* unresponsiveness as central to culpability and blame.⁷

The opposition thinks that rational or non-culpable mistaken belief or ignorance about what to do exculpates by excusing. Let's note three important things about this. First, the opposition needs to help us understand what it means for these defective beliefs to be rational or, at the very least, non-culpably held. Second, they need to help us see how such ignorance exculpates. Third, they need to help us understand how the relevant judgments could be mistaken.

It is important to remember that the relevant judgments are supposed to be *mistaken*. On some views about practical obligation and responsibility, there are epistemic constraints on practical norms or practical reasons such that the subject who follows her evidence or forms a justified belief about what to do might act in ways that are *bad* but could not act in ways that are *wrong*.⁸ Consider the increasingly popular idea that if you justifiably judge that you ought to X, it could not be that your

⁷ It is an interesting further question what role *de dicto* responsiveness should play in a theory of culpability. For all that we argue here a subject might be culpable for acting in the belief that her action is wrong.

⁸ See Zimmerman's (2008) discussion of the prospectivist view of obligation, for example.

obligation is to do other than X.⁹ If we think of the justified judgment as the one supported by the available evidence or the one that you can hold without violating any of your procedural epistemic obligations, this view seems to imply that the cases that we would describe as cases of wrongdoing would really be cases of justified action. Ignorance does not excuse, on this account, because it subverts obligation and thus functions as a crypto-justification.

We assume that these views about obligation are mistaken. The relevant epistemic factors *don't* subvert obligation, so if they have any significance for moral responsibility at all, it's because they help to excuse our failure to meet our obligations. The parties to our disagreement agree with this: obligations don't 'shift' as a consequence of our forming procedurally well-formed beliefs about what to do. Our disagreement is about the culpability we have for our actions when we do wrong while acting from moral ignorance.¹⁰

Our opponents often stress that the target mistaken moral judgments are rational, responsible, and non-culpably held. They seem to concede that it is harder to show that someone should be excused for acting on a culpable mistaken moral judgment. To clarify their conception of responsible belief it would be helpful to borrow an idea from Williams (1979). Think of a subject's subjective motivational set as something that's comprised of her beliefs, desires, and whatever further attitudes you think matter to the rationality of these beliefs and desires. Now think of the kinds of rational transitions involved in rational belief change. Let's say that a subject's procedural epistemic obligations are obligations to modify her attitudes in line with the kinds of transitions involved in rational belief change and let's say also that a belief is procedurally rational if produced by such transitions from a coherent starting set. On this way of thinking, we cannot say that a subject's moral judgment is culpable,

⁹ See Gibbons (2013) and Littlejohn (2012). They both accept a conditional like this one, but whereas Gibbons thinks that facts about your perspective determine what you can justifiably believe and do, we think that the reasons that bear on what to do are among the reasons that determine whether your beliefs about your obligations are justified.

¹⁰ We remind readers of these points because it's important, when working through the examples, that we don't lose sight of the fact that the relevant actions will be wrong and that the features that make the actions wrong are known to the relevant agents (albeit, not under that description).

irresponsible, or irrational simply because the initial elements of the subjective motivational set didn't include the right ones or included the wrong ones. The issue, then, is whether someone who has the 'bad' subjective motivational set can be culpable for her deeds even if she acts in line with beliefs about what to do that are procedurally epistemically rational. We think that that she can.

3. Zimmerman on Control and Blame

Zimmerman (1997, 2008) criticizes the Asymmetry Thesis on the grounds that it implies that we can be culpable for our actions even if we are non-culpably ignorant of the fact that we were obligated to do otherwise. His argument is designed to show that it is only possible to be culpable for an action if, at some point, the agent acts in the belief that her action is wrong. The argument goes as follows:

- (P1) Agnes did A and A was morally wrong but Agnes was ignorant of this moral fact simply because she didn't believe A-ing to be wrong.
- (P2) Agnes is culpable for A-ing only if she's culpable for the ignorance in which she A-ed.
- (C1) If Agnes was culpable for wrongfully A-ing, she must have been culpable for the ignorance in which she A-ed.
- (P3) Agnes is culpable for something only if she is in control of it.
- (C2) Agnes is culpable for wrongfully A-ing only if she's in direct or indirect control of her ignorance.
- (P4) But Agnes is not directly in control of her ignorance, so she could only be indirectly culpable for her ignorance.
- (C3) If Agnes was in control of her ignorance and culpable for her wrongfully A-ing, there must have been something else that she had direct control over that could have alleviated her ignorance.
- (P5) There was no such thing.
- (C4) Agnes was not culpable for having A-ed wrongfully in her ignorance.

The upshot is that all culpability can be traced to some case in which the subject acted wrongfully in the belief that her action was wrong. Knowledge of the wrong-making features is itself insufficient for culpability.

There is something to this argument. It does seem that we do not have direct control over whether we are ignorant of the relevant moral facts. It also seems that responsibility requires control. Nevertheless, it seems that a crucial premise in the argument is problematic.¹¹ A defense of (P2) is needed.¹² As a first stab, someone

¹¹ Here we point to some reasons why (P2) is unmotivated. In later sections, we will argue that (P2) is implausible on a certain (common) reading: i.e., if we assume that procedural rather than substantive considerations determine whether a subject is culpable for her ignorance or mistaken belief. Some authors seem to think that a certain kind of epistemic procedural excellence excuses the agent when she acts in light of her beliefs.

¹² A defense of (P2) is in order because it seems a controversial assumption in an argument against a well-established view about responsibility. An intuitive motivation for (P2), due to Rik Peels, might be: 'If Ayn engages in wrongdoing and is non-culpably ignorant of the fact that she wasn't permitted to A, how could we blame her? Don't we expect people to act from their propositional attitudes about what is right and what's wrong?' We don't think we do. First, it is not a general requirement for acting well that agents make decisions in light of the (undefeated) reasons that apply to them and determine what they should do. A similar point holds for verdicts. To borrow a point from Raz (1990: 181), it is no fault in Ayn that the reason she never killed someone is that she never had the slightest inclination to do so or that the thought of doing so never crossed her mind. So long as the thought doesn't cross her mind and she acts in acceptable ways, she does no wrong and we can find no fault with her, even though she might not be guided by thoughts about moral verdicts. Second, there is Ross' (1930) point about duties and motives. Ross attacked the idea that our duty to act in certain ways is invariably a duty to act in this way *from the motive of duty*, on the grounds that this further requirement conflicts with the categorical nature of the original duty. If moral duties are categorical, they are incumbent upon us whatever our motives happen to be. The duty to keep a promise is binding upon agents even when they're incapable of being moved by duty to keep the promise, e.g. when they cannot cultivate this particular motive. Something similar holds for attitudes. If we cannot settle the question whether we ought to do something, doing it might still be our obligation even if we couldn't be guided by the belief that it is obligatory. None of this amounts to an attack on (P2), but it explains why we should not invariably expect people to be guided by their beliefs about their obligations.

might offer a defense of this principle linking Agnes' culpability for A-ing to her culpability for A-ing while being in a certain state:

Act-State Culpability:

Agnes is culpable for A-ing when A-ing in state S only if she is culpable for S.

Everyone should agree that this is a silly principle. Agnes isn't culpable for being on planet Earth or being born after Moses' death, but it shouldn't follow from the fact that she's not culpable for those things that she's not culpable for A-ing. We need a more refined version of Act-State Culpability to get the argument moving.

Perhaps a better formulation would be this:

Act-State Culpability 2:

Agnes is culpable for A-ing when A-ing in state S only if she's culpable for S where S is a state that is part of a root requirement of responsibility.

If the relevant state were *not* such a requirement, then it would be hard to see how Act-State Culpability 2 could fare better than its predecessors. Unfortunately, if we do not have a defense of Act-State Culpability 3, this defense of the key premise seems to beg the very question at issue. We deny that a belief about the moral status of A-ing is a root requirement of responsibility.

Not only does Act-State Culpability 2 fail to provide independent support for (P2), it also seems to be mistaken. Consider three (putative) root requirements of responsibility: having control over whether you A; the fact that such and such a reason R is the reason that explains why it is wrong to A; being aware that it is wrong to A. Culpability does not require knowledge or belief that *these* conditions are met, so we cannot appeal to Act-State Culpability 2 to explain why culpability would require either culpable ignorance or awareness of the deontic status of the relevant action. All sides should probably agree (if only to avoid a potentially troubling regress) that if the agent freely decides to A in the knowledge that A-ing is wrong, she should be blamed for having A'd. Notice that an agent can A in the knowledge that A-ing is wrong without *knowing* that any of the three root requirements for responsibility just listed are satisfied, for example, without knowing that you are aware that A-ing is wrong.

Thus, according to Act-State Culpability 2, knowing that you are wrong to A is *not* sufficient for culpability because you would also need to know that you knew this and know that you did this freely. Thus, Act-State Culpability 2 is implausible.

We think that someone who thinks of responsibility as having to do with *de re* unresponsiveness would also be skeptical of (P2). They wouldn't see why it should matter to evaluating the quality of the agent's will whether she represents her actions as being impermissible. If it didn't matter to responsibility whether she represented her actions in this way, it is hard to see why it should matter to responsibility whether she had direct control over whether she believed or knew the moral status of her actions. We also think that people who think of responsibility as having to do with *de dicto* unresponsiveness should be skeptical of (P2). Suppose you were in the habit of writing down the day's tasks over breakfast. Your guardian angel decided to check your list one morning. She informs you that you really are obligated to do all but one of the ten things on your list and that you are obligated to refrain from doing one of the things on the list. In hearing this, you realize that your grounds for believing that you ought to do each of the things on the list are all equally strong and that your grounds for believing that you made a mistake are even stronger.¹³ In light of this, you suspend judgment on whether you ought to do the first thing, ought to do the second thing, ought to do the third thing, etc. so that you no longer believe of any item on the list that you are obligated to do that thing. Your angel helpfully offers to reveal to you what you are obligated to do and what you are obligated to refrain from doing, but you decline the offer and then complete the ten tasks you set for yourself.

Because you suspended judgment while completing your list, you did not believe of any particular action that it was prohibited. You could have easily remedied

¹³ The subject might still have some high degree of confidence that each task was required and some low degree of confidence that each task was prohibited. We note that the view in question states that the subject could only be blamed for completing the items on the list if either (a) the subject acted impermissibly in the belief that she was doing something impermissible or (b) the subject's ignorance was itself the result of a case of clear-eyed akrasia. Neither condition is met in this case because there is no reason to think that the agent's decision to decline the angel's offer was itself a case of clear-eyed akrasia. We agree that the subject is culpable for failing to let the angel remedy the agent's ignorance, but this is why we're skeptical of (P2) as Zimmerman understands it.

your ignorance but you decided not to.¹⁴ According to the line of argument just sketched, you could only be culpable for acting in this ignorance if you were culpable for the ignorance in which you acted. We would expect that the people who favor a view on which responsibility is a matter of *de dicto* unresponsiveness would say that you were culpable for your ignorance, but Zimmerman's argument rules out this response. You could only be culpable for your ignorance if the story contained something it did not. At some point in the story, you would have to knowingly engage in wrongdoing, but the story included no such belief.¹⁵ While we do not love views on

¹⁴ In response to this, Rik Peels reasonably asked why the agent would perform the actions given that the motivation was now gone. We would say that the suspended judgment removed only one potential motive for completing the day's tasks. It might be that there were further non-moral reasons for completing the list, so we shouldn't assume that the motivation for completing the list was gone, only that the agent didn't complete the list in the belief that the listed items were morally mandatory.

¹⁵ Guerrero (2007) argues (rightly, we think) that a problem for (P2) can be found by focusing on cases in which the subject is non-culpably uncertain about her obligation but culpable for her action when she acts in spite of suffering from this uncertainty. We think that there is a real problem for Zimmerman (and Rosen who seems to offer arguments nearly identical to Zimmerman's) precisely because the subject who declines the offer seems obviously culpable for her ignorance even though there is no point in the story in which the subject judges that she ought to take the available means for remedying her ignorance. Someone like Rosen might well say that she had an obligation to take the angel's offer but we see two problems on this front. First, even if Rosen is right about the obligation she has to remedy the ignorance, it does not follow that the failure to fulfill this obligation is culpable. Thus, if someone like Rosen wishes to explain why the purported counterexample fails by arguing that the subject is culpable for her ignorance, they need to tell a further story about why this is a *culpable* failure to fulfill an obligation and not merely a failure to fulfill an obligation. Couldn't the subject be excused for her failure? Second, it might seem outrageous to suggest that she should be excused for this failure, but this misses an important point. Like Zimmerman, Rosen thinks culpability for acting in ignorance requires that there is a point at which the subject acts impermissibly in the belief that they are acting impermissibly. Our agent forms no beliefs while believing that she should not and performs no actions while believing that she should not. If she is culpable for her beliefs, this has to be because some sort of *de re* unresponsiveness is the root of epistemic responsibility. We think that is plausible, but anyone who opts for this view of epistemic

which *de dicto* responsiveness is the key to responsibility, we can see the attraction of saying that someone would be culpable for turning down the information the guardian angel offered when she had no reason to turn it down.¹⁶ Such a decision, it would seem, would manifest a kind of *de dicto* unresponsiveness that seems to do a decent job of accounting for the intuition that the subject shirked her responsibility and should be blamed for it. It is probably worth noting that the intuition seems to get stronger the more mistakes the angel promises to point out. Fiddling with the numbers, however, does not introduce the thing that Zimmerman would need to vindicate the intuitions that the subject should be blamed for acting in her state of ignorance and should be blamed for failing to do the thing that would have remedied this ignorance.

As we see things, a crucial premise in Zimmerman's argument is unmotivated. We do not see any good reason to accept that Agnes' culpability for her actions turns on whether she would be culpable for the ignorance in which she acts. Even if Agnes is not culpable for failing to form correct beliefs about the moral significance of the facts she knows it still seems she might be culpable for failing to act as if these facts were morally significant.

4. Rosen on Fairness and Blame

responsibility has to explain where Zimmerman's argument goes wrong when applied to epistemic responsibility and culpable ignorance. If *de re* unresponsiveness is the key to understanding culpability for belief, their arguments must fail when applied to belief. In seeing that they fail, we might rightly be skeptical of their arguments concerning culpable action.

¹⁶ Rik Peels remarks, rightly, that Zimmerman would also have to say that the subject is blameless if she suspended judgment on whether she should turn down the information of the guardian angel and that seems less plausible. We happen to think that the best versions of the views that place *de dicto* responsiveness at the center of moral responsibility would have some requirement that agents take due diligence to determine what their obligation would be. Thus, we think that regardless of whether we take *de dicto* or *de re* responsiveness to be central we should agree that a subject can be culpable for the failure to improve her epistemic position even if this failure does not result from some case of clear-eyed akrasia.

Zimmerman's argument might rest on premises that are more controversial than needed to defend his view. Let's consider a slightly different rationale for that view. Consider Agnes just before she acts and then as she's acting. If she judges mistakenly that she should A and the mistake is based on a prior non-culpable mistake about the facts, we agree that it wouldn't be right to blame her. Suppose, however, that Agnes' mistaken moral belief about what she should do derives from a further mistaken moral belief. And suppose that she has met all of her procedural epistemic obligations in judging that she should A. Consider two questions:

The Epistemic Culpability Question: If we stipulate that Agnes has met her procedural epistemic obligations and so differs from someone whose moral judgments constitute knowledge simply because of a difference in their respective subjective motivational sets, couldn't Agnes be blameless in her mistaken moral belief?

The Practical Culpability Question: If we stipulate that Agnes' guiding moral belief is epistemically blameless, wouldn't she be morally blameless for acting on it?

Rosen believes that the answer to the Epistemic Culpability Question is obviously 'Yes'. Suppose that Agnes' guiding moral belief *is* epistemically blameless. Shouldn't we then say that she is morally blameless for doing what she did? Rosen thinks that the answer to the Practical Culpability Question is also 'Yes'.

Rosen thinks that when we distinguish these questions and think about them carefully focusing on the point at which Agnes has settled the moral question, we will have to agree that she is no more morally blameworthy for her action than she is epistemically blameworthy for her belief. Given that everyone seems to agree that she is blameless for her belief, he expects that we will now see that blame is governed by a fairness norm:

The Fair Blame Thesis: It is unfair to blame someone for doing something if he blamelessly believes that there is no compelling moral reason not to do it.
(2003: 74)

Unless there is some way for the subject to work out given the resources available to her that she shouldn't do what she judges she should do, we have to agree that moral ignorance and mistake exculpate. Without some procedural epistemic failure to point to, it is hard to see how we could blame Agnes for her beliefs. And if we think that she cannot be blamed for her beliefs, it does seem that we should not then blame her for putting those beliefs into action.

We worry, though, that we can get caught in two minds here precisely because we might try to answer the Epistemic Culpability Question while thinking about a broadly procedural notion of epistemic rationality and then try to answer the Practical Culpability Question by thinking about epistemic rationality in more substantive terms.¹⁷ Maybe we should be careful in answering the Culpability Epistemic Question, reflecting on whether our answer varies depending upon whether we are thinking of epistemic rationality in purely procedural terms. We should also think about whether our answer to the Practical Culpability Question might vary depending upon what notion of rationality we have in mind. Once we notice the danger of running these different conceptions of epistemic rationality together, we need to ask which of these notions of rationality is more closely connected to culpability.

We can see that these worries are not idle if we think about fairness and culpability in other domains. Let's consider two:

The Fair Legal Culpability Thesis: It is unfair to legally blame someone for doing something if he blamelessly believes that there is no compelling legal reason not to do it.

The Fair Epistemic Culpability Thesis: It is unfair to epistemically blame someone for believing something if he blamelessly believes that there is no compelling epistemic reason not to believe it.

In choosing these two, we have chosen sources of reasons and norms where, arguably, forming procedurally rational beliefs about how we should respond to the situation

¹⁷ For a discussion of the distinction see Hooker and Streumer (2004).

does not guarantee that we have conformed to these norms and met our obligations, in order to see whether it would be plausible to say that a procedurally rational belief about how we should respond would be exculpating.

The first thesis does not seem obviously true. If someone took due care and then broke the law in full awareness of the relevant facts because of ignorance or mistake about the law, it is not obvious that the law should give a full excuse. (Equally, it is unclear that a system would be unfair simply because it treated ignorance of law differently from ignorance of fact.) After all, it is important to use sanctions to encourage people to take extra steps to inform themselves about the law before acting. It is not obviously *unfair* to have a system that would not recognize mistake about the law or ignorance of the law as excuses. Someone could say in response that the law differs from morality in important ways, but we don't see anything inherently unfair in a social practice of holding people morally responsible in ways that deviate from the Fair Blame Thesis as a way of creating an incentive for people to take extra steps to see to it that their actions don't violate moral standards. Perhaps the Fair Blame Thesis is too weak and more forgiving than a practice that is not inherently unfair and that involves some attractive features a more forgiving practice would lack.¹⁸

The second thesis does not seem obviously true, either. Consider the relationship between (a) failures to meet procedural epistemic obligations and (b) the judgments about first-order failures that we would form if we met our procedural epistemic obligations while working from our subjective motivational set. Couldn't there be failures of mesh? We cannot think of any reason to suppose that the attitudes that we form in breach of our procedural epistemic obligations are ones that we could or would identify if we tried to identify them. And we cannot think of any reason to suppose that, if we tried to form attitudes about the first-order attitudes that we ought to have, the second-order judgments about our first-order attitudes would have to be reliably correct. Someone might blamelessly but falsely believe that there is no compelling reason not to believe *p* even when such a compelling reason exists and the subject is aware of it. (Of course, the subject would not be aware *that* this reason is

¹⁸ As David Owens and Rik Peels remind us, there are important differences between the law and morality that concern societal considerations about whether to blame, but our point is simply about the necessary conditions for it to be fair to blame.

such a reason.) It is not at all clear why the second-order belief excuses the first-order belief if the first-order belief is not supported by the evidence, held against the evidence, or resulted from some careless reasoning.¹⁹ Why would the presence of second-order belief shield the subject from blame or criticism on account of her unfortunate first-order beliefs?

The Fair Blame thesis might differ from these other theses, but it is not obvious that Rosen's Fair Blame Thesis can do the work he wants it to. To overcome initial skepticism, he needs to explain why we should think that being *epistemically* non-culpable in forming a belief ensures that the subject would be *morally* non-culpable for acting on this belief. The standards we use for assessing beliefs and actions differ. It is not implausible to think that the standards we use to determine culpability will differ depending upon the standard at issue. If Rosen can offer us something that could address this kind of skepticism, he will give us grounds for a kind of bridge principle. This principle can be used to argue in *two* directions. Any argument for the conditional claim that we cannot blame people for acting on blameless moral beliefs is, *inter alia*, an argument that we can blame people for forming moral beliefs that they would be blamed for acting on. Why shouldn't we turn the Fair Blame Thesis on its head? If Rosen is really right that anything that can excuse a mistaken moral belief is something that would excuse actions that manifest *de re* unresponsiveness, perhaps this is an indication that sometimes *nothing* excuses the action or belief.

¹⁹ To give but one example, consider the consistency norm for belief (i.e., a norm that says that we are required to avoid having inconsistent sets of beliefs) and an evidential norm for belief (i.e., a norm that says that we are required to have those beliefs that receive sufficient support from the evidence). It seems that these norms can pull us in different directions. It seems easy enough to imagine subjects who, in isolation, come to have different views about which of these norms is genuine. Perhaps Ayn could come to believe the consistency norm having met her procedural epistemic obligations and Agnes could come to believe the evidential norm having done the same. This, in turn, would result in situations where one of these subjects would be likely to form false beliefs about what she had reason to believe, but whichever norm is genuine should settle what reasons there were for these subjects to believe or refrain quite apart from their attitudes about which norm was genuine.

It should be noted that Rosen's account is very generous in terms of the excusing conditions it offers. The Fair Epistemic Culpability Thesis implies that we are almost never culpable in our beliefs. We can get at this result in two ways. First, he seems to agree with Zimmerman that ignorant wrongdoing is never a locus of original responsibility and the grounds for thinking this for action are grounds for thinking this for belief. We rarely, if ever, form bad beliefs as a result of some prior wrong that we acknowledged as wrong when we committed it. Second, if we apply Rosen's test within the epistemic domain, we will find that a subject's failure to discern her epistemic failings can render blameless the beliefs that would have to be culpably held for the agent to be morally culpable in her actions. A subject who reasons fallaciously (e.g., commits the gambler's fallacy or conjunction fallacy) will be expected to fail to see that there is a compelling epistemic reason for her to abandon her fallaciously formed beliefs. It is hard to see, given Rosen's restrictions, how we could hold people responsible for the beliefs they form by means of fallacious reasoning since many of the higher-order beliefs about epistemic reasons and requirements will be blamelessly held. If we think about Agnes who thinks that there is no compelling epistemic reason not to think that it's now much more likely that she will win after the next spin of the roulette wheel, it looks as if Rosen's account tells us that she cannot be blamed for her belief.²⁰

If we want to avoid giving out too many excuses for bad beliefs, we will want to drop the requirement that a subject can only be held responsible for her beliefs when there is a point at which she forms a belief she believes she should not hold. We will want to say that the failure to meet her procedural epistemic obligations, acknowledged or not, is something she can be held responsible for. And once we do that, we will have to opt for a kind of *de re* responsiveness approach to epistemic culpability. Once that is done, though, the question as to why we should be fixated on procedural epistemic obligations becomes pressing. If being held responsible for my beliefs can involve a failure to meet an epistemic obligation that is *not* some merely procedural epistemic

²⁰ Of course, Rosen can try to point to past cases of wrongful actions that result in the acquisition of intellectual vices or irrational beliefs, but it seems that this approach will never get the scope right: cases of clear-eyed akratic action that result in these epistemic imperfections are relatively rare.

obligation, why couldn't we say something similar about the responsibility for our actions?

5. A Remaining Challenge

We mentioned earlier that there were two very different positions about the relationship between normative ignorance and culpability. First, there is the view that normative ignorance or mistake remove or diminish moral culpability because they play a role in determining the moral obligations under which the subject finds herself. Second, there is the view that normative ignorance or mistake do not affect the moral obligations under which a subject finds herself but they can nonetheless remove or diminish culpability by exculpating the subject, and they do so unless the normative ignorance or mistake are the result of some culpable antecedent act of the agent's. We have examined two motivations for the second view and have found them unconvincing. Against those positions, there is the more controversial view, defended e.g. by Elizabeth Harman, that moral ignorance or mistake do not generally remove or diminish culpability:

We are morally obligated to believe the moral truths relevant to our actions (and thus not to believe false moral claims relevant to our actions), and we are often blameworthy for failing to meet these moral obligations, even if we have not been guilty of mismanagement of our beliefs, and even if our ignorance is not motivated.

Wrong actions that result from false moral beliefs are not thereby blameless; indeed, they may be loci of original responsibility. While both the beliefs and the actions are blameworthy, the actions are not blameworthy because the beliefs are blameworthy. Rather, the actions and the beliefs are blameworthy for similar reasons. (2011: 459)

We share Harman's skepticism that epistemic mistakes or ignorance about moral truths are typically blameless and hence exculpating. But here we want to focus on a presupposition that underlies arguments for our opponent's view, as well as a possible

objection to ours, that concerns rationality and epistemic culpability. The presupposition can be brought into view by reflecting on the objection. Consider the Enkratic Principle:

EP. If it is reasonable or rational to judge that you should X, it is reasonable or rational to X in accordance with your judgment.²¹

This principle tells us that there is a principled connection between the rational status of a belief about what to do and the rational status of the actions that the belief rationalizes. Once we settle that a conscientious person has reasonably judged that she should act in a certain way, writers like Rosen insist that we lose the grounds for blaming them for so acting. The principle itself does not tell us whether we should see the standards of rational or reasonable action as constraints on the standards for reasonable belief or whether we should use some independent conception of reasonable belief (e.g., belief that is formed as a result of meeting your procedural epistemic obligations) to tell us what it is to act reasonably. But many writers seem to think that the reason that this principled connection holds is precisely that the epistemic constrains the practical.²² As these philosophers see things, there is some independent conception of rational or reasonable belief, say, that it is supported by the evidence, and holding such a belief ensures that the things that it rationalizes (e.g., further beliefs, affective responses, decisions or actions) also count as reasonable. This is *not* because such beliefs infallibly direct us in such a way that we always hit some

²¹ See Broome (2013). Some of the principle's defenders assume that it should be understood in a top-down way. Foley (2001), for example, thinks that ordinary epistemic standards determine whether a belief is reasonable or rational and argues that certain actions count as reasonable or rational *because* the agent performed them in the belief that these actions were appropriate.

²² It can do so in two ways. Some use the principle to argue that a subject's perspective or evidence determines what she should do, defending views on which a subject's reasonable beliefs about what to do cannot lead them astray. Others use the principle more modestly and argue that a subject's perspective or evidence matters only to judgments about her responsibility for failing to meet practical obligations that are determined, in part, by matters that might be beyond the subject's ken. For discussions, see Foley (2001) and Gibbons (2013). An exception to the consensus is Arpaly (2000).

independent target, but because such targets turn out not to have any direct bearing on what is reasonable or rational to do. This ‘top-down’ approach is not required by the Enkratic Principle though it is popular among those who are attracted to it.

We find hints of this top-down approach in Rosen’s discussion. Rosen describes the case of Bill, who deliberates about whether to tell his wife a self-serving lie when ‘all things considered Bill should tell the truth and face the music’ (2004, 304). Bill knows that self-serving lies are morally wrong but ‘he has been badly taught ... through no fault of his own’ about the significance of moral considerations and, on the basis of his false beliefs about the importance of morality, he reaches the verdict: ‘All things considered, I should lie’. Rosen notes that Bill has been ‘as thoughtful and reflective and curious as a person should be in the circumstances’ (2004: 305). Surely, one might conclude, in these circumstances it is reasonable or rational for Bill to judge that he should lie and so, according to EP, it is therefore reasonable or rational for him to lie.

This seems to land us into trouble. For in holding that, contrary to what Rosen says, Bill would be blameworthy for telling the lie in those circumstances, we are committed to the seemingly unfortunate view that an agent can legitimately be morally criticized for doing what was rational or reasonable for him to do (see e.g. Rosen, 2003: 74-75). While it is unclear what precisely people mean when they say that someone acted reasonably, a plausible interpretation is that they either acted with justification or should not be blamed for the failure to do so. Either way, it seems that the judgment that Bill’s actions were reasonable seems in tension with the view that he’s not morally excused.

The conclusion that our view has such an undesirable consequence, however, depends on the assumption that Bill’s judgment about what to do was rational. But why should we think this? This is not something the Enkratic Principle forces us to say. If we accepted that principle, we only have to say that reasonable beliefs about what to do and reasonable actions ‘sway together’. Those who insist that Bill’s actions were reasonable are relying on the presupposition about rationality mentioned above, namely that rationality is always and everywhere a matter of procedure. But, couldn’t Bill’s failure to discern the weight of competing reasons or perceive the force of some fact itself be a kind of irrationality? Couldn’t that have some bearing on the epistemic rationality of Bill’s normative judgment about whether to lie?

Our response is two-fold. First, the objection assumes both the Enkratic Principle and the top-down approach. If we reject one or both of these, the objection loses its force. We think that the worst view to take here is one that combines the Enkratic Principle with the top-down approach. Let's suppose the Enkratic Principle is true. We see no good reason to prefer a view on which the epistemic merits of Bill's belief, such as they are, shield him from criticism for his deeds to a view on which we appeal to the normative standards that Bill tries to get right in criticizing him for both his actions and the attitudes that rationalized them. It may be helpful here to think about emotion and another of Rosen's baddies. Smith is Rosen's sexist American father of the 1950s. Smith treats his sons and daughters differently, encouraging his sons to pursue their education, let's say, but not his daughters 'because they're girls' (2003: 67). Suppose his daughters protest the unequal treatment, and leave home go to University and Smith gets upset in turn. This kind of emotional response might well mesh with his outlook: he may be upset because he mistakenly thinks that his daughters are doing a terrible job doing what they should do. What should we say about his emotional response to the situation? The things that upset him are the very things that show that his daughters' lives are going well and which make them happy and their lives meaningful. We don't think that it's a datum that his emotional response is reasonable, nor that we shouldn't blame him for his reprehensible emotional responses to his daughters' flourishing. Either there is no principle like the Enkratic Principle that links the reasonableness of emotional responses to judgments about how one should feel about things, or there is such a principle but the very things that make Smith's emotional responses unreasonable threaten to show that his moral appraisals are equally unreasonable.

Second, our suggestion is that, even if theoretical rationality is essentially no more than procedural rationality, practical rationality is not.²³ Failures of practical rationality are not limited to failures to be thoughtful, reflective and curious about one's moral beliefs and values or about what follows from them: they include also failures to be aware of the relevant moral facts or properly to appreciate their weight in

²³ If we were forced to say that theoretical rationality is essentially a matter of procedural rationality, this response would force us to abandon the Enkratic Principle, but we see no good reason to think that it is.

various circumstances. Rosen disagrees and says that, in his example, Bill's wife ought to judge that his lie 'was a perfect manifestation of practical rationality' because 'through no fault of his own he found himself believing that all things considered, he should lie' since he believed 'that in the circumstances, his selfish interests were more important than my [his wife's] moral interests' (2004: 306).

We maintain that Bill's action was not a perfect manifestation of practical rationality. It was a rather imperfect manifestation of such rationality because in acting as he did he didn't act well. To be sure, he acted in a way that was consistent with his values and with the judgements derived from those values, so that he was procedurally practically rational, both in his judgement about what to do and in his action. But his judgement was mistaken because he failed to appreciate that what he had been taught about the importance of morality was wrong. That failure led him to conclude that he ought to place his selfish interests before his wife's and that was a failure to appreciate the relevant moral facts, that in this case constituted a failure of substantive practical rationality.

This conception of substantive practical rationality has its roots in Aristotle and it has been defended recently by, among others, Warren Quinn and Philippa Foot. Quinn conceives of practical reason as having as its 'primary job the correct evaluation of ends, actions, and qualities as good and bad in themselves... To the extent that one realizes or serves some such good one acts well. To the extent that one realizes or serves some such bad one acts badly' (1993: 30). And he adds:

On my view, the only proper ground for claiming that a quality is rational to have or an action rational to do is that the quality or action is, on the whole, good. (1993: 46)

Accordingly, one does not exemplify perfect substantive practical rationality if, like Bill, one is mistaken in one's evaluation of whether in doing something – say tell a self-serving lie – one would be acting well.

This may seem an excessively substantive and ambitious conception of practical rationality and one may question why we should accept it. That challenge cannot be properly addressed here but we think that it can be partly met, or at least and

indication of how it could be met can be provided, by reminding ourselves of Philippa Foot's rendering of Quinn's central insight:

We think of practical rationality as a kind of master virtue. But what would be so important about it if it would license or even mandate disgraceful actions? How in such a case could it keep its supposed status as the master virtue? Why should its criteria have to be met in a satisfactory theory of moral virtue?
(2004: 5)

In other words, if practical rationality is not action-guiding in this substantive (moral) way, why should we care much about being practically rational? One might respond that we should care about being practically rational because it is the most promising way of realising our most cherished projects, values and desires. There are two things to say about this response. One is that that may indeed be a reason why we should care about being practically rational but it need not be the only one. The other is that, our most cherished projects and values may be themselves irrational or unreasonable and, as it has been plausibly argued, if we have no good reasons to pursue our projects and values then we cannot have good reasons to do what would lead to their realisation, so mere procedural practical rationality cannot sustain our concern to be practically rational.²⁴

Our view seemed to have the undesirable consequence that an agent can legitimately be criticized for doing what was rational or reasonable for him to do. Our response is that the objection depends on a conception of practical rationality (mere procedural rationality) that we regard as too limited to do all the work that the notion of practical rationality can plausibly be called to do. And, given this richer conception of practical rationality, it is not obvious that our view implies that an agent can legitimately be criticised for doing what was practically rational for him to do. It does imply that an agent can legitimately be criticised for doing what was merely procedurally practically rational for him to do not but we regard that as a defect of our account.

²⁴ For discussion of these issues, see Quinn (1993), Korsgaard (1997) and Raz (2005).

There are two further possible objections to our view that we can envisage. One is that our position violates the principle that moral responsibility requires control. For it seems that we don't have control over whether we have the required moral knowledge, sensitivity, acumen, insight, etc., and yet we claim that we are typically blameworthy for our normative mistakes or ignorance and for the wrong actions that flow from them. We deny that we lack the requisite control over whether we have the required moral knowledge.²⁵ In general, most adult human beings have the capacity and opportunity for moral reflection to achieve the required moral knowledge to guide them in most of their actions (and, as we said above and Harman stresses, they also have an obligation to do so).²⁶ Whatever Bill was taught as a child about the importance of one's selfish interests, if we presume that Bill is a normal human adult, he will have reached the intellectual maturity required to be able to realise that it is morally wrong to, as a matter of principle, place one's selfish interests over the legitimate interests of others.²⁷

The second objection relates to the remarks about Bill's ability. What does our position say about someone who simply cannot know the relevant moral facts, or appreciate their relative importance in different circumstances? After all, we must admit that some people are very poorly equipped to do so, if at all, perhaps because they were exposed to extreme forms of abuse and brutality that destroyed or thoroughly distorted their capacity for moral thinking,²⁸ or because they have limited

²⁵ Some writers might distinguish control, whether voluntary or not, from influence here and insist that we have indirect influence over our beliefs, but not control. We think such influence is all that is required.

²⁶ The same is not true for non-moral facts: we don't always have the ability or obligation to know all the non-moral facts that determine whether, in the particular circumstances in which we find ourselves, phi-ing would amount to wrongdoing.

²⁷ So we worry about the precise statement of control conditions for reasons discussed above in connection to Zimmerman's argument. There are some cases in which it seems that we should all agree that an agent's culpability for acting in ignorance does not trace back to some case of clear-eyed akrasia.

²⁸ We are assuming extreme cases that impair one's capacity for moral thinking and people who just receive poor or defective moral education are may or may not fall into this category. The suggestion is consistent with the thought that people with higher and better attuned capacities,

intellectual capacities, or because of ‘an undetected brain anomaly or the like—some blameless condition’ (Rosen 2004, 305). Our response is that those people are not exculpated from their moral obligations by their moral ignorance; rather, they are not subjects of the obligations in the first place.

6. Conclusion

Ignorance and mistaken belief are often perfectly good excuses. It seems plausible to think that they are so whenever the ignorance or the mistaken belief is non-culpable. The plausibility of the thought evaporates, however, if we get the standards for culpability for ignorance or mistaken beliefs wrong. We have argued that we get them wrong if we think, as Zimmerman and Rosen do, that culpability requires an instance of known wrongdoing: that acting wrongly requires *de dicto* unresponsiveness to one’s obligations. We have also argued that this mistake may be encouraged by linking the question about culpability for epistemic imperfections to an anaemic notion of practical rationality (i.e., one that reduces practical rationality to some procedural notion), both in judgement and action. We have argued that there is no reason to think that a plausible answer about culpability and excuse for wrongdoing can be guided by claims about rationality that rely on that anaemic conception of practical rationality. So we hope to have shown that ignorance and mistaken belief are indeed often perfectly good excuses – but far less often than our opponents claim.²⁹

References

whether this is a natural gift, the result of superior education, or whatever, have stricter moral obligations and hence the bar for exculpation for them may be higher than for others.

²⁹ We would like to thank Oliver Black, Julien Dutant, David Owens, Rik Peels, the editors of this collection, an audience at King’s College London, and anonymous referees for discussion and feedback on earlier versions of this paper.

- Arpaly, N. 2000. On Acting Rationally against One's Best Judgment. *Ethics* 110: 488-513.
- 2001. *Unprincipled Virtue. An Inquiry into Moral Agency*. Oxford University Press.
- Broome, J. 2013. *Rationality through Reasoning*. Wiley.
- Foley, R. 2001. The Foundational Role of Epistemology in a General Theory of Rationality. In L. Zagzebski and A. Fairweather (ed.), *Virtue Epistemology*. Oxford University Press, pp. 214-30.
- Foot, P. 2004. Rationality and Goodness. *Royal Institute of Philosophy Supplement* 54: 1-13.
- Gardner, J. 2007. The Gist of Excuses. In his *Offences and Defences*. Oxford University Press, pp. 125-40.
- Gibbons, J. 2013. *The Norm of Belief*. Oxford University Press.
- Guerrero, A. 2007. Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philosophical Studies* 136: 59-97.
- Harman, E. 2011. Does Moral Ignorance Exculpate? *Ratio* 24: 443-68.
- Hooker, B. and Streumer, B. 2004. Procedural and Substantive Practical Rationality. In *The Oxford Handbook of Rationality*. P. Rawling and Alfred R. Mele (eds). Oxford University Press, pp. 57-74.
- Korsgaard, C. 1997. The Normativity of Instrumental Reason. In *Ethics And Practical Reason*, Cullity and Gaut (eds), Oxford University Press.
- Littlejohn, C. 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- Quinn, W. 1993. Putting Rationality in its Place. In his *Morality and Action*, Cambridge University Press. Reprinted in: *Value, Welfare, and Morality*. R. G. Frey and C. W. Morris (eds.). Cambridge University Press, pp. 26-50.
- Raz, J. 1990. *Practical Reason and Norms*, 2nd Edition. Oxford University Press.
- 2005. The Myth of Instrumental Rationality. *Journal of Ethics and Social Philosophy* 1: 1-28.
- Rosen, G. 2003. Culpability and Ignorance. *Proceedings of the Aristotelian Society*, 103: 61-84.
- 2004. Skepticism about Moral Responsibility. *Philosophical Perspectives* 18: 295-313.

- Ross, W.D. 1930. *The Right and the Good*. Oxford University Press.
- Wieland, J.W. Forthcoming. What's Special about Moral Ignorance? *Ratio*.
- Williams, B. 1979. Internal and External Reasons. Reprinted in: *Moral Luck*,
Cambridge University Press, 1981, pp. 101-13.
- Zimmerman, M.J. 1997. Moral Responsibility and Ignorance. *Ethics* 107: 410-26.
- 2008. *Living with Uncertainty*. Cambridge University Press.