Responsibility – The Epistemic Dimension II
Abstracts

George Sher
Blame and Moral Ignorance
Can a person legitimately be blamed for acting wrongly when he knows what he is doing, but doesn't know that it is wrong? Like a good many others, I believe the answer is sometimes "yes," but that whether blame is appropriate in any given case depends on certain facts about the agent's epistemic situation. My aims in this paper are to establish, first, that a morally ignorant wrongdoer's epistemic circumstances do have a bearing on his culpability, but, second, that giving content to this familiar view is far harder than is generally appreciated.

Nomy Arpaly
No pre-read

Alexander Guerrero
Difficulty and Responsibility
In many cases of ignorance, we appear to have done nothing wrong, we have only been unlucky: unlucky in what we were taught, unlucky in the evidence we happen to encounter, unlucky in that our usually reliable method has in this case led us astray. Consider the example of the Ancient Slaveholder, who holds slaves, but at a time when (we are to suppose) no one – neither slaveholders nor slaves – questioned the morality of slavery; it was just taken for granted as a morally acceptable cultural practice. Many suggest that he has just been unlucky in this way. The claim made on the Ancient Slaveholder's behalf is not that he is unlucky in that it would have been *impossible* for him to see that slavery was wrong. That slavery is wrong is an a priori matter, and there were a priori arguments that were available – in some sense – even back then. Instead, the claim made on the Ancient Slaveholder's behalf is that he is unlucky in that it would be *exceptionally difficult* for him to come to believe that slavery was wrong. Gideon Rosen says that "[g]iven the intellectual and cultural resources available to a second millennium Hittite lord, it would have taken a moral genius to see through the wrongness of chattel slavery."

Here, then, is a natural suggestion. Difficulty: If coming to believe some proposition *p* is exceptionally difficult for an agent A through no fault of A's own, then A is not morally responsible for failing to believe *p*. As I discuss in the paper, there are variants of this principle that concern action as well. This suggestion is in tension with recent work by Pamela Hieronymi, who has argued that "[w]e are fundamentally responsible for a thing … because it reveals our take on the world and our place within it – it reveals what we find true or valuable or important." Hieronymi (and Arpaly and Harman) would reject Difficulty. Others, such as Rosen, would accept it, and argue that it excuses the Ancient Slaveholder. In this paper, I argue that we should accept Difficulty, but deny that it excuses the Ancient Slaveholder.

I begin by distinguishing several different senses of difficulty: probabilistic, intellectual/epistemic, and physical/ergonomic; and several different places of potential intellectual/epistemic difficulty: identifying that there is a question, believing the correct answer, and having that belief play an appropriate role in one's theoretical and/or practical reasoning.

I argue that Difficulty and nearby variants of Difficulty can be offered with all of these understandings, and that many of these are plausible given our ordinary concept of moral responsibility. I also argue that since difficulty comes in degrees, excuse by way of Difficulty will also come in degrees, and so must be reformulated.

I argue, using a number of different examples of responsibility for non-moral belief and responsibility for physical actions, that Difficulty is a central part of our concept of moral responsibility for beliefs and actions, and that there is no plausible, principled reason for denying this with respect to Difficulty with respect to moral belief.

I conclude, however, by arguing that we must be careful how we understand and apply Difficulty. I argue that, on plausible understandings of Difficulty, it is *not* difficult for the Ancient Slaveholder to hold the correct moral belief, and so the Ancient Slaveholder is either not excused at all or only weakly or very partially excused. I conclude with some thoughts about the generality of excuse through Difficulty.

Elinor Mason

The Epistemic Condition and Quality of Will

A very plausible account of moral responsibility connects responsibility with the quality of an agent's will. This raises the question, what about an agent's will is important? Is it necessary that an agent's will is transparent to her in order for to be properly blameworthy or praiseworthy for what she does? I argue that we should distinguish between two sorts of bad will, and correspondingly, two sorts of blameworthiness. Ordinary blameworthiness, requires moral knowledge: an agent must be aware that she has a bad will to be blameworthy in this way. The other kind of blameworthiness, objective blameworthiness, applies to agents who do not understand their own will. I argue that it makes no difference whether their ignorance is through lack of capacity or through external circumstances. Agents who do not, as a matter of fact, understand the moral quality of their will are outside of our moral community in an important sense, and in a different category of appraisal.

Elizabeth Harman

When is Failure to Realize Something Exculpatory?

Sometimes one has sufficient evidence for p, but one simply doesn't realize that p is true. In some of these cases, the failure to realize is exculpatory. For example, if some vital medicine is in a safe, and opening the safe would require solving a very difficult math problem, one is not blameworthy for falling to solve the problem (nor for failing to open the safe), if one tried as hard as one could to solve it. One isn't blameworthy even if the solution to the math problem was not fundamentally out of one's reach. In other cases, I've argued, a failure to realize something is not exculpatory. In general, failures to realize the moral truth are not exculpatory. In a third kind of case, I will argue, failures to realize something are not exculpatory, though they might seem to be exculpatory. Consider someone who believes he is simply flirting with a woman, and asking her out, but is really harassing her. If he does not realize he is harassing her, this is not exculpatory, I claim. If he does not realize how much he is bothering her, this is not exculpatory, I claim. My claims raise a puzzle. Why is failure to solve the math problem exculpatory, while failure to realize what is really going on in the harassment case is not exculpatory? I will lay out and address this puzzle.

Endre Begby

Can Moral Responsibility Exceed Epistemic Responsibility? Applying the Distinction between Justifications and Excuses

There is a quite general moral principle which tells us that we ought to take every reasonable precaution to avoid harming innocent people. This principle applies readily to our actions, but also – though perhaps in a more indirect fashion – to our beliefs. Prejudiced belief might serve as a case in point: prejudice can cause substantial harms to innocent parties, both in terms of emotional damage and in terms of the effects of social exclusion.

The no-harm principle accordingly has an epistemic dimension: it is tempting to assume that this moral notion of responsibility-for-belief must track our normative assessments of people's epistemic standing more generally. Accordingly, for instance, Nomy Arpaly and Miranda Fricker have argued that prejudiced belief necessarily involves some form of epistemic culpability, some specifiable mishandling of evidence. Recent work, however, throws doubt on this assumption: many cases of prejudiced belief may well turn out to be in good epistemic standing (cf. Begby 2013). This finding leaves social epistemology in an uncomfortable position: if those who harbor prejudices aren't *culpable*, then we would struggle to make sense of the idea that those who face prejudices are *wronged* – i.e., unjustifiably harmed – as a result. And if they aren't wronged, why should society at large be concerned to mitigate the disadvantages they nonetheless face as a result?

This paper seeks to develop the idea that our moral responsibilities can reach beyond our epistemic responsibilities: where one's beliefs concern other people (or stand to affect social relations more broadly) one can be morally responsible to *the truth* in a way that exceeds one's epistemic responsibility to *the evidence*. To make this point I draw on the legal distinction between justifications and excuses: an action that harms another can be justified only if it is based on a *true* belief. Actions that manifest prejudiced beliefs are clearly not justified in this sense, since prejudices are by definition false. They may, however, be excusable, insofar as they are based on a blameless assessment of the available evidence. Nonetheless, even an excused agent can be held responsible for such harms: in the legal setting, epistemic blamelessness may exempt one from criminal prosecution, but not from compensating for unjustified harms. This argument, if successful, would allow us to maintain that prejudice results in a distinctive kind of unjustified harm even when the prejudice itself is epistemically blameless.

Errol Lord

On The Intellectual Conditions for Responsibility: Acting for the Right Reasons, Conceptualization, and Credit

There are at least two faces of responsibility. On the one hand, the actions for which we are responsible can count against us in a certain way. They can be actions for which we are blameworthy. On the other hand, the actions for which we are responsible can count in our favor. They can be actions for which we are creditworthy.

In this paper I am primarily interested in the question of which intellectual conditions must be met in order to act in a creditworthy way. What I say about this will directly tie up to a broader theory of the intellectual conditions for responsibility more generally. I focus on credit because I am primarily interested in the prospects of a particular theory of what it is to act in a creditworthy way. According to this theory, The Right Reasons Theory, what it is to act in a creditworthy way is to act for the right reasons, i.e., act for the reasons that make the act performed right.

In order for the Right Reasons Theory to be plausible, I take it, it needs to turn out that agents with ordinary intellectual powers can act for the right reasons. This is because most ordinary people have a modicum of virtue and thus act in a creditworthy way at least some of the time. I take this as a given. At the very least, failing to explain this claim is a mark against a theory of creditworthiness.

This paper has a negative aim and a positive aim. The negative aim is to show that the kind of Right Reasons Theory defended by Arpaly & Schroeder (2014) has a very hard time securing the claim that people with ordinary intellectual capacities act for the right reasons. This is surprising since one of the main motivations for their theory turns on cases involving characters acting in a creditworthy way even though they seem to lack a certain intellectual pedigree. I argue that their view fails because it imposes a *Conceptualization Condition* on acting for the right reasons. According to this condition, in order to act for the right reasons, one must have mental states whose contents are conceptualized via some normatively relevant concepts. I will argue that this puts too high an intellectual burden on acting for the right reasons and thus puts too high an intellectual burden on acting in a creditworthy way. I will argue further that the failure of this view makes the prospects of any conceptualization condition look dim.

That's the negative aim. The positive aim is to use the failure of the Conceptualization Condition to motivate a different type of theory of acting for the right reasons. The aspect of this theory that is relevant here is that acting for the reasons essentially involves the manifestation of a certain kind of *know how*. In particular, in order to act for the right reasons, one has to manifest knowledge about how to use those reasons in particular ways. I will argue that this theory can explain the data that motivates the Conceptualization Condition without demanding that ordinary people have exceptional intellectual powers. This theory thus avoids the problems that plague Right Reasons Theories that posit a Conceptualization Condition.

Gunnar Björnsson

Explaining (Away) the Epistemic Condition on Moral Responsibility (II)

Ignorance excuses, but *when* and *how* is controversial. This paper argues that "the epistemic condition" on moral responsibility is best understood as resulting from a more basic requirement, the *explanation* condition.

Generally, an agent is to blame for something (a decision, action, or outcome) only if it is suitably connected to relevant "motivational structures"—structures guiding the agent's responses to reasons. Correspondingly, moral blame always partly targets the agent's failure to properly take into account available considerations, i.e. some morally imperfect structures.

The relevant connection between motivational structure and object of blame, I argue, is that the structure is what I call a "normal significant explanation" of the object of responsibility. When there is *no* explanatory connection, there is no responsibility: the agent has no control over the object. When the explanatory connection is *not significant* (in a sense explained in the paper), there is not responsibility. And when there is an *abnormal* connection (again in a sense explained in the paper), there is again no responsibility.

This explanation condition on responsibility can account for and thus be supported by a number of responsibility phenomena not having to do with the agent's knowledge or lack thereof, including phenomena traditionally understood to follow from the control condition on moral responsibility. At the same time, it straightforwardly

accounts for paradigmatic cases where ignorance excuses, understanding these as cases where ignorance breaks the explanatory connection. Moreover, unlike strong epistemic conditions, it allows for cases where ignorance does not seem to excuse, and unlike prominent weaker conditions (e.g. FEC in Sher's *Who Knew?*), it provides an elegant account of these cases. These are simply cases where imperfect motivational structures normally and significantly explain failures to pay attention to ethically significant information and resulting failures to take ethically important action or ensure ethically important outcomes.

Gwen Bradford
Hard To Know
Is the epistemic condition for moral responsibility sensitive to difficulty? If it would be hard to know whether an action is morally wrong, does this shape responsibility for that action? It's a natural thought that, all else equal, if discerning some morally relevant factor would be exceptionally difficult, we are not to blame if we fail to recognize it. Restricting the discussion to factual ignorance, this paper argues that difficulty per se does not shape the epistemic condition or any other aspect of moral responsibility. To see this, we need to understand difficulty. According to the best account of difficulty, difficulty is a matter of exerting effort. The same act may involve a lot of effort from one agent, but not from another. There is no stock set of what we may call effort-requiring features. Rather, what matters for difficulty is relative to the agent. So long as an act would take a significant amount of effort for an agent, it is difficult for them. Equipped with this understanding, we can see that it is not the difficulty per se that mitigates blame in cases where it does, but the effort-requiring features that are relevant for responsibility. Interestingly, this holds both for the epistemic condition of moral responsibility, and for any aspect of the action itself. If discerning morally relevant features of an action would be difficult for Anna because she is developmentally disabled, her limited cognitive abilities mitigate her responsibility. In contrast, if discerning morally relevant features of an action would be equally difficult for Bob, who is not disabled but only extremely lazy, Bob is fully responsible and blameworthy if he fails to muster the effort. In other cases where it appears difficulty is relevant, and effort-requiring features are not relevant for responsibility, difficulty may be extrinsically significant, insofar as it shapes the moral valence of the action. As a result in these cases difficulty is not relevant for moral responsibility, but for the rightness or wrongness of the action, and for this reason it shapes praise- or blameworthiness. And this too holds for the epistemic condition or any other aspect of the act itself.

Justin Caouette
Do Psychopaths Meet the Epistemic Condition on Moral Responsibility?
Many have claimed that psychopaths are not morally responsible because they fail either the epistemic condition or the control condition. In this essay, I would like to focus on those arguing for exculpating the psychopath on grounds that they fail to meet the epistemic condition. Two philosophers: Ish Haji (2010) and Neil Levy (2007, 2007b), have recently argued that psychopaths are not morally responsible because they fail to meet the epistemic condition (but for different reasons). I explore the arguments employed by Haji and Levy and argue that they are not sufficient to exculpate the psychopath. Haji claims (1) impairment in moral perception excuses the psychopath of moral responsibility and (2) impairment in moral perception moderates the degree to which he is culpable, I argue that we do not have sufficient evidence to

warrant (1) and that (2) could be correct, however, I point to some reasons to be weary of degree of moral responsibility talk. The same can be said of Neil Levy's (2007, 2007b) argument in favor of morally exculpating psychopaths because they lack "the relevant moral knowledge for distinctively moral responsibility". To argue against both Levy and Haji I suggest an epistemic condition that stays true to our normal ascriptions of praise and blame yet allows for reasonable mitigating circumstances to circumvent those ascriptions. Certain impairments do inhibit one from attaining moral knowledge; however I argue that psychopathy is not one of them. Psychopathy causes the agent to lack some ethical sensitivity but ethical sensitivity is not as paramount as Haji claims nor does it leave the psychopath with no moral knowledge as Levy asserts. Thus, if psychopaths are not morally responsible for their actions it must be because they fail a condition other than the epistemic one.

Matt King
Responsibility Without Tracing
I consider two ways in which tracing might apply to the epistemic condition. First, it might be that certain cases can only be explained via tracing. Second, it might be that the condition can only be satisfied when the epistemic relations have the right sort of history. I argue that tracing does not apply in either of these ways. On the first score, one might think that tracing is required to explain responsibility in cases of culpable ignorance. Ordinarily, ignorance excuses, but not when the person is themselves at fault for their ignorance. I argue instead that ordinary explanatory mechanisms of action can explain responsibility in these cases without tracing. On the second score, one might think that epistemic states that could not be properly 'traced' would undermine responsibility for subsequent action (e.g., because the state were implanted or otherwise abnormally generated). I argue, in contrast, that the relevant epistemic relations are importantly ahistorical, and thus tracing is again unnecessary.

Matthew Talbert
Akrasia, Awareness, and Blameworthiness
Some philosophers believe that because ignorance tends to excuse, a blameworthy actor will either be a knowing wrongdoer or her ignorance will be traceable to a prior instance of knowing wrongdoing. I argue that this claim is false because a wrongdoer's behavior can exhibit moral and attributional qualities that are sufficient for blameworthiness even if she never acted akratically. However, while blameworthiness does not depend on a wrongdoer recognizing the moral status of her behavior, it often does depend on the wrongdoer being aware of the likely consequences of her behavior. Without this latter sort of awareness, an agent's behavior may lack the moral features that make blaming responses appropriate. Moreover, what is often required in this context is actual awareness of the consequences of one's behavior. Thus, the mere fact that an agent could have or should have been aware of the consequences of her behavior is in many cases not sufficient for genuine moral blameworthiness.

Milo Nanni
Akrasia, Culpable Ignorance, and the Risk of Moral Impermissibility
A morally ignorant agent does not know that he should do, or refrain from doing, a certain action. This may be due to factual ignorance or to ignorance of some moral rule. Contention surrounds when wrong actions done from moral ignorance are blameworthy. I argue that they are blameworthy only when moral ignorance is culpable

and that, therefore, moral ignorance excuses when it is non-culpable. I defend three claims from Rosen:

(1) A wrong action done from ignorance (WI) is blameworthy only if the agent is blameworthy for that ignorance.

(2) Ignorance is blameworthy only when it results from the violation of procedural obligations in the management of one's beliefs.

(3) The mismanagement of belief out of which culpable ignorance arises involves akrasia.

Guerrero and Harman attempt to show the falsity of (1). Harman also rejects (2) and (3). Guerrero and Harman argue for an account that assesses actions as blameworthy in a larger number of cases than (1)-(3) allow.

Guerrero and Harman reject (1) by arguing that an agent can be blameworthy for WI even if he is not blameworthy for that ignorance. I show that this argument fails because they mislocate the relevant state of ignorance.

Harman suggests that cases of *affected ignorance* – in which the agent chooses not to be informed of what he should know – do not result from the violation of procedural obligations in the management of one's beliefs: this would constitute an objection to (2). I show that affected ignorance is culpable only if the agent chooses not to be informed of something *he believes* he should know, because of its moral relevance.

In rejecting (3) Harman looks to cases in which culpable moral ignorance can be said to derive from the non-akratic exercise of vices. I argue that even if this derivation were correct, it would entail that the agent should act against his best judgement, a claim I find intuitively wrong.

Miriam McCormick

When Do False Beliefs Excuse?

When thinking about what agents must know to be morally responsible, or when lack of knowledge can excuse, what is often of concern is some kind of moral knowledge, or an appreciation of the nature of the actions one is performing, e.g. that a particular act will result in harming another and that is morally wrong. It has been suggested that severe cognitive incapacities and some mental illnesses excuse one from moral responsibility because they render one incapable of appreciating moral facts. For example, many have argued that psychopaths cannot be moral responsible for such reasons (e.g. Levy 2007).

In this paper I am not primarily concerned with when lacking *moral* knowledge or having false *moral* beliefs excuses one from moral responsibility. I am interested, rather, in how we should think about people in very poor epistemic circumstances, cases where access to evidence is limited or controlled. Now if someone has a very distorted picture of reality, this will likely also likely lead to faulty moral beliefs as well, but I think sometimes the issue in cases where it is unclear whether or not the "epistemic condition" is met, it is because we think the person who lacked some knowledge or awareness *should not* have lacked such knowledge or awareness. (Sher 2009) And so the question about responsibility gets pushed from responsibility for the action to responsibility for the ignorance or lack of awareness.

It is important then to think about epistemic or doxastic responsibility in trying to clarify the epistemic condition on moral responsibility. I will consider three cases, one where we are inclined to excuse (or severely mitigate) responsibility, one where we are not, and one where I think it is harder to assess. I will argue that reflection on

these cases show that many recent accounts of responsibility for attitudes (including beliefs) that attempt to divorce responsibility from control (what I call “character-based” and “reasons-responsive” accounts) are flawed because the key to understanding the difference in the cases requires thinking about what control the agents have or lack.

Paulina Sliwa

The Justified Belief Principle

My paper argues that the epistemic condition on moral responsibility involves norms of inquiry. The first part of the paper argues that agents are not fully morally responsible for actions which they justifiably believe to be right – even when those actions are in fact wrong. Call this the Justified Belief Principle (JBP). I give two arguments in support of JBP. The first appeals to cases of misleading moral testimony and advice. When agents do the wrong thing because they are misled by someone they reasonably regard as a reliable moral advisor, we do not hold them fully morally responsible for their wrongdoing. Such cases, then, lend intuitive support to JBP. Second, I offer a more general argument in support of JBP. I argue that an agent is not fully morally responsible for acting wrongly, when she does the best she can to do what’s right. An agent who has a justified belief that her action is right, is doing the best she can. I defend this argument against several objections.

JBP tells us when agents are not fully morally responsible for their actions. But it also bears on which epistemic condition an agent must meet in order to be fully morally responsible for their action. The second part of this paper explores these implications. Specifically, insofar as the JBP is correct, a plausible epistemic condition on moral responsibility should give us some explanation for why it holds. I consider the following candidate epistemic conditions:

> EC: An agent is fully morally responsible for an action only if she is in a position to know its moral status.

I argue that EC gives us a satisfying explanation of why JBP is true and that it’s also independently attractive. I end by raising a challenge to EC: how do we spell out, in a principled way, when an agent is in a position to know the moral status of her action? I suggest that, surprisingly, one way to address the challenge is by moving to an even stronger epistemic condition on moral responsibility:

> EC*: An agent is fully morally responsible for an action only if she knows its moral status.