



## Doing One's Reasonable Best: What Moral Responsibility Requires

*ABSTRACT: Moral responsibility, I argue, requires agents to do what is within their abilities to act morally. This means that an agent is to blame just in case his wrongdoing is due to an underperformance, that is, to a failure to do what he can to act morally. I defend this account by considering a skeptical argument about responsibility put forth by Gideon Rosen and by Michael Zimmerman. I explain why the epistemic condition they endorse is inadequate and why my alternative epistemic condition, which directly follows from my general condition on culpability, should be preferred. I then defend my view against potential criticisms.*

**KEYWORDS:** ethics, responsibility, metaphysics, free will, agency, action

A person who freely and knowingly acts wrongly is to blame for her wrongdoing. But can an agent be culpable for a wrong action she does not believe is wrong? In my view, she can. Agents who act wrongly because they fail to perform according to their abilities, I will argue, are culpable for their wrongdoings. Hence, an agent is to blame if her bad act is based on a false belief she ought not to have formed, given her cognitive abilities. My goal in this paper is to flesh out this position.

I will start by examining a skeptical argument about responsibility put forth by Gideon Rosen and by Michael Zimmerman. The value of this exercise is twofold. First, the skeptical argument relies on a widely accepted epistemic condition on culpability. Showing the inadequacy of this condition will help motivate my alternative epistemic condition. Second, Rosen and Zimmerman provide arguments in support of this condition. Explaining where these arguments go wrong and why the conception of responsibility they rely on is misguided will enable me to shed light on the requirements of moral responsibility. I will show how my proposed epistemic condition naturally follows from these requirements. The rest of the paper will address possible objections to my view.

I would like to thank Sherri Irvin and the anonymous referees for this journal for their very helpful suggestions about this essay. I am also grateful to Neal Judisch, Jay Kelly, Matt Priselac, and Wayne Riggs for useful conversations about the topics discussed here.



## 1. Skepticism about Responsibility

There is general agreement that an agent may be blameless for her wrongdoing. Two kinds of excuse are deemed admissible. First, the agent may lack the relevant control over her action. The second kind of excuse is epistemic: the agent may not have realized that her action was wrong. Traditionally, skeptical arguments about moral responsibility have concerned the first type of excuse: if our behaviors are governed by (deterministic) laws of nature, then, the arguments go, we do not have genuine control over these behaviors. But I will examine a different kind of skeptical argument, put forth by Gideon Rosen (2004) and Michael Zimmerman (1997, 2008), that exploits the second kind of excuse.

For example, Dr. Wong prescribes a common antibiotic to her patient. Unbeknownst to her, the patient is allergic to the antibiotic and is harmed by taking it. Is Dr. Wong culpable for prescribing the antibiotic? It depends on the details of the story. Suppose first that the patient's allergy was noted on his chart, which Dr. Wong failed to consult. In this case, Dr. Wong is to blame for her action. (More carefully, Dr. Wong is *indirectly* to blame for her action if she is *directly* culpable for her failure to consult the chart. I will examine the distinction between direct and indirect culpability shortly.) But now suppose that Dr. Wong took reasonable precautions: she consulted the patient's chart, which did not mention the allergy; she asked the patient to confirm the relevant entries on the chart, and so on. In such a case, Dr. Wong is not blameworthy for the wrong prescription. Cases such as this one are invoked to support the following thesis:

- (1) If agent S performs action A from ignorance, then S is culpable for the act only if S is culpable for the ignorance from which she acts.  
(Rosen 2004: 300; Zimmerman 1997: 414).

What does 'ignorance' amount to exactly? Suppose a rather invasive biopsy is advisable when and only when blood tests come back positive for a certain antigen. One day, Dr. Dimon performs the biopsy on a patient whose blood test is negative. He does that to collect the health insurance payment. However, unbeknownst to Dr. Dimon, the lab technician that day was a disgruntled employee who filled out the forms randomly. As a matter of luck, the result he recorded on this particular patient's form was accurate. Hence, although he truly believes that performing the biopsy is morally impermissible, Dr. Dimon does not *know* that it is. Yet, he is blameworthy for performing it (see Rosen [2008: 596–97] for similar examples). Ignorance should thus be equated with a lack of true belief rather than a lack of knowledge. Finally, both Rosen and Zimmerman hold that if the agent has a *dispositional* rather than *occurrent* belief that her action is morally wrong, she counts as ignorant that her action is morally wrong (I will examine their reasons for holding this in section 4.2). Thus, (1) amounts to:

- (1') If agent S performs action A while lacking the occurrent true belief that her doing A is morally wrong, then she is culpable for doing A only if she is culpable for lacking that occurrent belief.

This brings us to the phenomenon of *culpable ignorance*. Morality, Rosen points out, makes epistemic demands on us. He writes:

As you move through the world you are required to take certain steps to inform yourself about matters that might bear on the permissibility of your conduct. You are obliged to keep your eyes on the road while driving, to seek advice before launching a war and to think seriously about the advice you're given; to see to it that dangerous substances are clearly labeled, and so on. These obligations are your *procedural epistemic obligations*. (2004: 301)

And as Rosen makes clear, 'The procedural obligation is not itself an obligation to know or believe this or that. It is an obligation to *take steps* to ensure that when the time comes to act, one will know what one ought to know' (2004: 301).

Procedural epistemic obligations, Rosen points out, are highly dependent on one's situation and impossible to codify. Several factors must be taken into consideration when determining an agent's procedural epistemic obligations. One must consider how likely it is that new evidence will affect the ranking of the available courses of action. In particular, one needs to consider how probable it is that, in light of the new information, one's current best option will no longer be at the top. One also needs to consider how morally superior an alternative course of action would be, relative to the current favorite, in light of the new evidence that would be gathered. Finally, one needs to consider the costs involved in acquiring the new information and how these compare to the costs of acting wrongly and to the benefits of a right action (Jackson 1986).

Two points emerge. First, a procedural epistemic obligation is—despite its name—a *moral* obligation. A failure to fulfill one's procedural epistemic obligations is thus a moral failure. And like any moral failure, it may or may not be culpable. Second, strictly speaking, a culpably ignorant agent is not culpable for her ignorance (i.e., lack of true belief), but for a past failure to fulfill her procedural epistemic obligations. Culpable ignorance is thus always *indirect* or *derivative*. In other words,

- (2) An agent S is culpable for the ignorance from which she acts only if her ignorance is the upshot of some prior culpable act or omission.

That is, (1) and (2) entail that an agent who culpably acts from ignorance must be derivatively blameworthy for her action. Her culpability must be due to some earlier *original*, or *direct*, culpable failure to fulfill a procedural epistemic obligation. This means that an agent who is directly culpable for her action cannot be acting from ignorance. Direct culpability is thus associated with the following epistemic condition:

- (EC1) An agent S is directly blameworthy for her wrongdoing A only if S has an occurrent belief that her doing A is morally wrong.

On the proposed view, direct culpability is due to what Rosen calls *clear-eyed akrasia*, which occurs when an agent acts while having an occurrent true belief that she is doing something wrong. EC<sub>1</sub> is required for Rosen and Zimmerman's argument to hold up. Suppose EC<sub>1</sub> is false and that there are cases in which an agent is directly blameworthy for acting wrongly from ignorance, that is, while lacking an occurrent belief that she is acting wrongly (I will describe several such cases in the following sections). Then, contrary to what (1) and (2) entail, it is not the case that an agent can only be derivatively blameworthy for acting from ignorance.

Let us take stock. Culpable wrongdoing is either clear-eyed or from ignorance. According to (1) and (2), the latter must be traceable to an earlier culpable wrongdoing. But to avoid a regress, that wrongdoing must be clear-eyed. We thus have:

- (3) An agent S is culpable for her wrongdoing A only if her doing A is, or derives from, an episode of clear-eyed akrasia (Rosen 2004: 307; Zimmerman 1997: 418).

Rosen and Zimmerman draw slightly different conclusions from this argument. Rosen distinguishes akratic actions from instances of what he calls *ordinary weakness of the will*. He writes: 'The akratic agent judges that A is the thing to do, and then does something else, retaining his original judgment undiminished. The ordinary moral weakling, by contrast, may initially judge that A is the thing to do, but when the time comes to act, loses confidence in this judgment and ultimately persuades himself (or finds himself persuaded) that the preferred alternative is at least as reasonable' (2004: 309). Akratic actions, Rosen contends, are extremely hard to distinguish from actions resulting from ordinary weakness of the will, even from the first-person perspective. The problem is that cases of ordinary weakness of the will are not the locus of original responsibility since they are instances of wrongdoing from ignorance. Rosen thus concludes that 'it would be unreasonable to repose much confidence in any particular positive judgment of responsibility' (2004: 308). Zimmerman (1997: 425; 2008: 176), on the other hand, concludes that because cases in which ignorant behavior can be traced to an episode of direct blameworthy action are rare, many ordinary apparently blameworthy actions are in fact blameless. Zimmerman's conclusion is thus stronger than Rosen's. Rosen concedes that the conditions for culpability may frequently be satisfied, but only concludes that confident judgments of culpability are never justified.

## 2. Contesting the Argument

A crucial premise of the skeptical argument is EC<sub>1</sub>. It is worth noting that epistemic conditions on culpability similar to EC<sub>1</sub> are widely shared. As a matter of fact, many authors endorse an even stricter condition that requires *knowledge* of wrongdoing. In her classic paper on culpable ignorance, Holly Smith (1983) calls an act done from ignorance an *unwitting act* and a violation of a procedural epistemic obligation

a *benighting act*. She writes: 'To say the culpably ignorant agent is to blame for his unwitting act is to say nothing more than that he was culpable in performing the benighting act, that it gave rise to the unwitting act, and that he knew at the earlier time that he risked this outcome' (1983: 566). More recently, McKenna (2012: 15) holds that blameworthiness for an action requires *knowledge* that the action is wrong (see also Fischer and Ravizza 1998; Ginet 2000; Haji 1997; and Levy 2005 and 2011 for similar epistemic conditions).

But EC1 is implausible. Suppose agent S is unsure whether her doing A is morally correct. Let us say that S attaches 0.5 credence to the proposition that doing A is morally wrong. S would be blameworthy for (freely) doing A. As Elizabeth Harman writes, 'If someone acts wrongly while genuinely unsure whether her action is wrong, we need not investigate whether she is blameworthy for being unsure to know whether she is blameworthy: she may well be blameworthy simply for doing what she did, which she believed might well be wrong' (2011: 449). Hence, culpability does not require the belief that one's action is morally wrong. To avoid Harman's counterexample, one may propose the following condition:

(EC2) An agent S is directly blameworthy for her wrongdoing A only if S lacks the belief that her doing A is morally permissible.

Surely, to avoid culpability, a person need not be actively thinking that she is doing the right thing. Hence, according to EC2, a person is to blame for her act only if she lacks either a dispositional or an occurrent belief that her act is permissible. But even if EC2 is interpreted this way, Harman's counterexample does not quite support it. We are ignorant of countless morally relevant facts and for this reason have no flat-out beliefs about the morality of many of our actions (and omissions). In many cases, we have at best a relatively high level of confidence that we are acting permissibly. It seems harsh to hold that we are blameworthy for every wrong action we perform under these states of mind. Suppose you must purchase a certain item and have the choice between two brands, A and B. (Let us assume, to simplify, that not buying is not an option.) The prices, apparent quality, and so on of A and B are the same. You are aware that all other things being equal, it is preferable to buy merchandise from manufacturers that have less morally objectionable practices. But through no fault of yours, let us suppose, you lack any specific information about the makers of A and B. You decide to buy A. Although you have some confidence that buying A is morally acceptable, you lack a flat-out belief that it is. Unfortunately, the manufacturer of A is a morally corrupt company that should not be patronized. You are not to blame for your purchase, it seems. If, on the other hand, your credence that buying A is morally acceptable were *lower* than your credence that buying B is, then you would be to blame for your purchase. Hence, instead of EC2 we should have: an agent S is directly blameworthy for her wrongdoing A only if S lacks the belief that her doing A is morally permissible *and* the credence S attaches to the moral permissibility of A is not at least as high as the credence S attaches to the moral permissibility of any other available course of action. However, talk of credence can be extremely cumbersome. For this reason, I will stick to the less accurate but more convenient talk of belief.

But EC<sub>2</sub> is problematic for other reasons. Consider an agent S who has fulfilled all her relevant procedural epistemic obligations. However, S underperforms, epistemically speaking: S's belief about the morality of her action is not well supported by her evidence. Because of that, she ends up (freely) performing a morally wrong action. S is blameworthy for her wrongdoing, even though she believes she is doing the right thing.

Like most doctors, Dr. Singh commonly prescribes antibiotics to patients suffering from acute bronchitis. This is a mistake, according to experts, since acute bronchitis is almost always caused by viruses. Moreover, antibiotics are often harmful, because they may cause side effects and contribute to antibiotic resistance. Upon hearing that at a medical conference, Dr. Singh promises herself that she will change her practice. A few days later, a patient with acute bronchitis comes for a consultation. Unfortunately, Dr. Singh fails to recall the promise she made to herself and prescribes an antibiotic. Although she believes that she is doing the right thing, Dr. Singh is to blame for prescribing an antibiotic. I am assuming that this is a lapse: Dr. Singh does not suffer from any memory deficiency or cognitive malfunction. A few hours later, let us suppose, when she suddenly realizes her mistake, she feels embarrassed. She had no excuse for doing what she did and would appropriately blame herself for her wrong prescription.

We commonly blame people for mistakes like that one. For example, a (usually reliable) bank teller is to blame for entering the wrong balance due a customer because of a miscalculation; a (competent) ambulance driver is blameworthy for taking a wrong turn and losing precious time during an emergency call if her mistake is due to her misreading the map. In each case, we may suppose that the agent sufficiently investigated the situation and thus did not violate any procedural epistemic obligation. EC<sub>2</sub> should thus be replaced by:

- (EC<sub>3</sub>) An agent S is directly blameworthy for her wrongdoing A only if S lacks an epistemically reasonable belief that her doing A is morally permissible. (Again, bear in mind that to avoid blame the agent need not have an occurrent belief. Moreover, a more accurate condition should incorporate a clause concerning reasonable credences.)

Why does moral responsibility require an epistemically *reasonable* rather than a *justified* belief? For example, Dr. Garcia faces a particularly difficult case. Although he has performed an array of tests on his patient, he has been unable to figure out what she is suffering from. What should he do next? He could perform more tests, but which ones? He could start a treatment, but which one? At this point, let us suppose, the available data are so complex that a correct synthesis would require access to a sophisticated computer software. The best Dr. Garcia can do, given his limited time, resources, and cognitive abilities, is rely on heuristic methods that often, but not always, yield the right course of action. Suppose his belief about what to do next is reasonable, given his current limitations, but that the data actually support a better course of action. Dr. Garcia's belief about what to do is reasonable but not, strictly speaking, justified. It would be uncharitable to blame Dr. Garcia

for his suboptimal recommendation. This is why moral responsibility requires not a justified but a reasonable belief that one's action is permissible. And what counts as a reasonable belief is agent-relative, that is, it depends on the agent's relevant cognitive abilities and background knowledge.

This means that agents with limited cognitive abilities, such as children and cognitively challenged people, may still have reasonable beliefs about the morality of their actions. A child's belief may appear unreasonable to more mature agents, but it counts as reasonable in the relevant sense here, provided that it is supported well enough, given the child's limited cognitive capacities. But what about seriously impaired agents, whose acts are based on completely irrational beliefs? For example, Bea suffers from Capgras syndrome and believes that people close to her have all been replaced by identical-looking impostors. Surely Bea is not to blame for acting on that belief. Yet, it seems odd to say that Bea's belief is *reasonable*. However, on the proposed account, Bea is *exempt* from blame, not because her beliefs about the identity of her interlocutors are reasonable, but because she lacks the ability to form accurate beliefs about them. As I will explain in the next section, according to the norm of moral responsibility, an agent must act morally *to the extent of her abilities*, including her abilities to form accurate beliefs about her surroundings. Hence, a being who lacks the relevant abilities is simply not subject to that norm: since she does not meet the metaphysical condition for morally responsible agency, at least in that domain, she is exempt from blame (or praise for that matter).

### 3. Doing One's Reasonable Best

To support condition EC<sub>3</sub>, I have invoked the fact that we do blame people for their wrongdoings when the latter result from cognitive underperformances. This condition, I have argued, is well supported by many of our everyday judgments. But further support can be gained from examining cases of culpability other than moral. Intuitions about these are perhaps less likely to be corrupted by one's theoretical commitments.

For instance, a chess player fails to notice that moving his pawn will allow his opponent's bishop to capture his rook by a so-called 'long diagonal'. This is a common mistake. But players are not all assessed in the same way for making it. While we readily excuse the novice for such a blunder, we blame the experienced player for it. We would say things like 'He, of all people, *should not* have moved that pawn,' or 'Someone like him *ought* to have anticipated the threat.' This suggests that two different norms are in place here. Both the novice and the experienced player violate a *primary* norm, according to which one ought to avoid *long-diagonal blunders*—that is, roughly, one ought to prevent one's pieces from being captured by a long diagonal, unless that allows one to make a bigger capture.<sup>1</sup> But there is a *secondary* norm that only the experienced player violates in this case. This

<sup>1</sup> Unlike moral norms, this chess-playing norm is *hypothetical*: if one's aim is to win (or at least play well), then one should play a blunder-free game. But this difference does not affect my point, since the chess players we are considering do want to play well and win.

secondary norm, which is dependent on the primary norm, is a norm of responsible chess playing. Responsible chess playing does not require the perfect satisfaction of primary norms. It requires, rather, that one play according to one's capacities. The experienced chess player violates this secondary norm because he was perfectly able to avoid the long-diagonal blunder. He is to blame because he underperformed relative to his abilities.

According to the secondary norm violated by the experienced player, a player S is to blame for a particular move M only if (roughly) S lacks a reasonable belief that M is not a long-diagonal blunder. (Note, once again, that to avoid blame a player need not have an *occurrent* belief that he is not making a long-diagonal blunder; a dispositional belief suffices.) Since the experienced chess player possesses a well-entrenched, reliable ability to detect the threat posed by his opponent's bishop, it is perfectly appropriate to blame him for his blunder. By contrast, the novice lacks a secure and robust ability to 'see the whole board'. He is thus not subject to this secondary norm and exempt from blame for making the long-diagonal blunder. Of course, he may well be to blame for not avoiding more obvious blunders that beginners can spot relatively easily.

In general, we respond differently to a person who fails to perform a desirable action because she lacks the required capacity than we do to a person who possesses that capacity but underperforms. The boundary between the two is vague, but there are clear instances of each category. For example, we blame the professional basketball player, but not the beginner, for missing an easy layup. Unlike the beginner, the professional is to blame, because she fails to perform according to her motor skills.

This is how, in my view, we should think of moral responsibility. We have a primary moral obligation to act according to what morality requires. And moral responsibility imposes a *secondary* moral norm on us. To act in a morally responsible way is to do one's reasonable best to respect one's primary moral obligations. And to do one's reasonable best is to perform according to one's relevant abilities, which include cognitive, volitional, and motor abilities.

The examples I gave in the previous section involved agents who lack epistemically reasonable beliefs. Omissions form another common type of underperformance. Randolph Clarke (2014: 164) describes a scenario in which his wife calls him at the end of the workday, asking him to buy milk on his way back home. Between his office and the store, he starts thinking about the paper he is working on and continues to think about it until he arrives home, where he suddenly realizes that he has forgotten to get the milk. Clarke writes that his omission is grounds for blame. Furthermore, he remarks, his culpable omission is not due to a prior blameworthy action. There are various things he could have done to ensure that he would not forget to stop and get the milk. However, he notes, these would have bordered on compulsion, given that he usually has no trouble remembering his plans. Clarke concludes, correctly in my view, that he is to blame for not getting the milk because his lack of awareness of his obligation to get the milk 'falls below a cognitive standard that applies to [him], given [his] cognitive and volitional abilities and the situation [he] is in' (2014: 167; see also Harman 2011: 463; Sher 2009: 24; Smith 2005: 236; and Smith 1983: 545 for similar examples).



Other instances of underperformance concern volitional capacities. We blame the person who knowingly fails to fulfill a moral obligation if it is clear that she possesses all the requisite abilities, including volitional ones, to do the right thing. However, we deem the sorrowing or depressed agent much less culpable for the same wrongdoing. A natural explanation of this difference is that a person's volitional capacities are somewhat compromised by sorrow or depression. Consider a father who snaps at his unruly child. We would find him less worthy of blame if we learned that he was already under considerable duress from some unrelated matter. Again, this is because we understand that hardship may negatively impact one's capacity for self-control. We would not exonerate that same behavior if it occurred in normal circumstances. In such a case, the father's culpability would be due to an underperformance: he would fail to exercise his capacity for self-control successfully.

What are capacities? I use the words 'abilities' and 'capacities' interchangeably, and in my view, the most promising analyses of capacities appeal to counterfactuals. Michael Smith offers a plausible account: 'Capacities are essentially general or multi-track in nature, and ... therefore manifest themselves not in single possibilities, but in whole rafts of possibilities' (2003: 27; see also Fara 2008; Fischer and Ravizza 1998; and Vihvelin 2004). For example, an agent has the capacity to answer a certain question correctly if he has 'the capacity to think of the answer to a whole host of slight variations on the question that he was asked, variations in the manner in which the question was asked, and perhaps in the exact contents of the question, and in the time of the question, and so on' (2003: 27). Furthermore, his correct responses in the nearby possible worlds must be explainable by a certain underlying internal structure. As Smith notes, an agent who, like Block's (1981) Blockhead, contains a set of internal mechanisms, each of which is dedicated to giving a specific response to a specific question, is not intelligent and should thus not count as having the relevant rational capacities. When these conditions are satisfied, we can truly say that the agent could have thought of the relevant response even if he did not, 'in one perfectly ordinary sense of "could"', as Smith (2003: 20) puts it.

Three points are worth making. First, does the ability to do A require that there be a possible world that is identical in history and laws to the actual world, except for the fact that the agent does something other than A? Like Smith (2003: 21), I am inclined to say 'no'. But incompatibilist readers should feel free to supplement the present account with this condition since none of the points I make here assumes the compatibility of determinism with the kind of abilities that moral responsibility requires.

Second, it is customary to distinguish between an agent's *general* abilities, that is, what she is able to do in a large range of circumstances, and her *specific* abilities, that is, what she is able to do now, in some particular circumstances (Mele 2003). Consider again the ambulance driver. Let us assume that in training, she has excellent map reading abilities. However, she has poor stress tolerance. As a result, she tends to misread maps when responding to actual emergencies. What should we say about her ability to read maps during real emergencies? One could insist that in these circumstances, she still has the (general) ability to read

maps efficiently, but that stress *masks* the manifestation of that ability. But one could also hold that she lacks the specific *ability-to-read-maps-efficiently-during-real-emergencies*<sup>2</sup> since in nearby worlds in which she faces real emergencies she is a rather inefficient map reader. Which notion of ability we adopt matters to some debates (see, for example, Whittle [2010] for a useful discussion of how this question affects our assessment of Frankfurt cases and the Principle of Alternate Possibilities). Fortunately, it is not necessary to settle this question here. This is because on each view, the ambulance driver is blameless for her mistake. On the ‘general-ability’ view, the driver is *excused* from blame because her (general) ability to read maps efficiently is defeated by stress, whereas on the ‘specific-ability’ view, she is *exempt* from blame because she lacks the (specific) ability-to-read-maps-efficiently-during-real-emergencies. Hence, my account of blameworthiness does not force us to choose among various levels of ability.

Third, consider an agent S with a general ability to do A, who is in favorable circumstance C. By assumption, then, S also has the specific ability-to-do-A-in-circumstance-C. On my understanding of abilities, S may attempt to do A in C without succeeding. In other words, although having an ability demands a measure of robustness and control, it does not entail having a perfect record (in favorable circumstances). J. L. Austin expresses this point well: ‘Consider the case where I miss a very short putt and kick myself because I could have holed it. It is not that I should have holed it if I had tried: I did try, and missed. It is not that I should have holed it if conditions had been different: that might of course be so, but I am talking about conditions as they precisely were, and asserting that I could have holed it. There is the rub’ (1970: 218). We find it appropriate for Austin to blame himself, or to “kick himself” as he puts it, for missing the putt—assuming, of course, that he does not have an inflated opinion of his golfing abilities. The point is that a well-intended agent with the relevant abilities may try but fail to do the right thing. In such a case, she is directly culpable for her blunder since it is within her abilities to do the right thing.

Moral responsibility thus places a heavier burden on agents than what Rosen and Zimmerman hold. Their mistake is to assume that a failure to exercise one’s capacities successfully is blameworthy only when it is deliberate, that is, only when the agent has an occurrent belief that she is acting wrongly. This assumption, I have shown, clashes with our everyday judgments. We blame Dr. Singh for failing to apply the knowledge she gained a few days earlier at a medical conference. And our judgment is not contingent on whether her failure was clear-eyed. Her failure is culpable because, we are assuming, she was clearly capable of avoiding it. Assessments such as this one, I have argued, are regularly expressed. We blame people for their blatant errors in judgment. The world of sports provides us with many illustrations of this phenomenon. Coaches and referees are blamed for their bad calls, especially when they are experienced and deemed competent enough to have avoided them. Similarly, players are castigated for missing easy plays.

Two key points emerge. First, having an ability does not guarantee success, even in favorable circumstances. We grant that in some cases, the circumstances

<sup>2</sup> The hyphens indicate that the circumstances mentioned are part of the ability in question.

may actually be *unfavorable* for the successful exercise of the person's abilities (the person was sick, grieving, under medication, etc.). But we do not hold that a failure to carry out an attempt to do A automatically entails that the circumstances were unfavorable for the manifestation of the person's ability—or that the person lacks the ability-to-do-A-in-those-circumstances. This is not how we think of abilities. Second, people are blamed for failing to perform according to their abilities, even when their failure is unintentional.

## 4. Objections and Responses

Rosen and Zimmerman might protest that it is pointless to object that their argument goes against commonsense since they openly admit that it does. My objections, they may add, amount to a mere appeal to ordinary judgments that they explicitly put into question. But this does not accurately represent the dialectic. First, I have shown that the cases of direct culpability that Rosen and Zimmerman describe can be accounted for by a weaker epistemic requirement EC<sub>3</sub>, rather than by their strong condition EC<sub>1</sub>, a key premise of their argument. Second, I have presented many counterexamples against EC<sub>1</sub>. These show that, contra Rosen and Zimmerman, we do not base our attributions of blame on the assumption that a (directly) culpable agent must be knowingly acting badly. I will now go further and examine their (and other authors') arguments in support of EC<sub>1</sub> and show that they are inadequate. A successful skeptical argument may well have a counterintuitive conclusion, but it ought to be based on well-supported premises. In this section, I will argue that Rosen and Zimmerman's argument does not meet that desideratum. I will proceed by considering objections to my view and offering my responses.

### 4.1 Rationality and Responsibility

**Objection:** In your view, agents are to blame when they act on the genuine but unreasonable belief that their action is permissible. Imposing such a burden on agents is wrongheaded. Moral responsibility demands that an agent perform an action only if performing that action is something the agent can do rationally. But what an agent can do rationally is a function of what she takes to be the case rather than what a better informed person would take to be the case. It is clearly irrational for an agent to act based on a belief she does not have. This is why EC<sub>1</sub> is the correct epistemic condition: to be culpable for her wrongdoing, an agent must see herself as acting wrongly.<sup>3</sup>

**Response:** This objection assumes an implausibly weak constraint on rational action. Consider again the blunder performed by the experienced chess player. It

<sup>3</sup> See Levy (2011: 128) for this line of thought. Levy objects to FitzPatrick's (2008) response to Rosen's argument. Like me, FitzPatrick denies that culpability requires clear-eyed akrasia. According to him, culpability may result from the voluntary exercise of vices such as arrogance, laziness, and dogmatism. I do not have the space to explore how FitzPatrick's virtue-theoretic approach compares to the deontological view I propose here. See Robichaud (2014) for criticisms of Levy's response to FitzPatrick.

makes perfect sense to hold that although he does not believe that he is making a blunder, he does not act rationally. Contrary to what the objection states, rationality does not merely demand that one act based on what one believes; it demands that one act based on the *evidence* one has. Since (by assumption) the chess player saw the position of his opponent's bishop but failed to take it into consideration, moving his pawn the way he does is not the rational thing to do.

Although it goes too far, the objection is correct in one respect: the demands of moral responsibility should take into consideration the agent's *perspective*. The unfortunate inhabitant of a world run by a Cartesian evil demon may act rationally, in the sense relevant to moral responsibility. She should not be blamed for bad actions resulting from a failure to detect the demon's machinations. This is why the notion of evidence relevant for responsibility should be construed in an *internalist* manner: since my twin who inhabits the demon world and I have the same (internalist) evidence about the morality of our actions, our beliefs on such matters are equally rational. (This does not preclude the possibility that externalist notions of evidence and rationality are better suited for other philosophical purposes.) Now, contrary to what the objection assumes, an agent's perspective includes not only what she believes, but also her evidence. Consider the bank teller again: clearly, given her evidence, her belief about the balance due the customer and her action based on that belief are not rational.

## 4.2 Culpability and Occurrent Belief about Wrongdoing

Objection: Zimmerman does offer support for his epistemic condition

(EC<sub>1</sub>) An agent S is directly blameworthy for her wrongdoing A only if S has an occurrent belief that her doing A is morally wrong.

He writes:

[I]f a belief is not occurrent, then one cannot act either with the intention to heed the belief or with the intention not to heed it; if one has no such intention, then one cannot act either deliberately on or deliberately despite the belief; if this is so, then the belief plays no role in the reason for which one performs one's action; and, I am inclined to think, one incurs culpability for one's action only if one's belief concerning wrongdoing plays a role in the reason for which one performs the action. (2008: 191)<sup>4</sup>

Response: Zimmerman's claim, according to which direct culpability requires that one act either deliberately on or deliberately despite the belief that one is doing the wrong thing, is implausible. A police officer grows frustrated at a demonstrator who refuses to heed his order to disperse. The police officer is irked by the demonstrator's

<sup>4</sup> Zimmerman allows for one possible exception: 'It may be that routine or habitual actions are performed for reasons to which one does not advert' (2008: 191). In such cases, he adds, one may be directly culpable for one's wrongdoing.

inaction, which he considers lawless, and decides to punish him by repeatedly striking him with his baton. The thought that he is doing something wrong never enters his mind and thus plays no role in the reason for which he performs his action. He simply acts out of the desire to hurt the demonstrator and the belief that the best way to do so is to bludgeon him. Surely, the fact that the police officer has no occurrent belief about wrongdoing does not excuse him from blame. This case provides us with a counterexample to EC<sub>1</sub>. But it does more: it shows that Zimmerman's defense of EC<sub>1</sub> is inadequate. Contrary to what he says, one may be directly culpable for one's action even though the belief that one is acting wrongly plays no role in the reason for which one acts. The police officer is directly culpable for his act, because he ought to and easily could have realized that his act was wrong. The condition according to which one may be to blame for acting while lacking a reasonable belief that one's act is morally acceptable better captures the way we think about culpability.

One might protest that the police officer's culpability is *indirect* here: what he is *directly* culpable for is his failure to stop and think about the morality of what he is about to do. This may be correct, but it merely displaces the problem. The police officer did not have any occurrent belief that it was wrong not to stop and think. He is directly culpable for his omission since he ought to and easily could have realized that he needed to consider the morality of what he was about to do.

Perhaps EC<sub>1</sub> could be salvaged by appealing to a slightly different understanding of ignorance. According to Rosen, 'X does A from ignorance when X acts in ignorance of every wrong-making feature of his act' (2008: 594). Note that a wrong-making feature of an act is not merely a feature that *contributes* to the wrongness of an act, but a feature in virtue of which the act is wrong. Hence, an act that possesses a wrong-making feature is necessarily a morally wrong act.<sup>5</sup> Rosen's epistemic condition is thus

(EC<sub>1</sub>') An agent S is directly blameworthy for her wrongdoing A only if S is aware of at least one wrong-making feature of A.

EC<sub>1</sub>' entails that the police officer is directly culpable for his act, for he is (plausibly) aware of a wrong-making feature of his act, namely, the harm he causes to the demonstrator. But EC<sub>1</sub>' does not fare as well with respect to agents such as Dr. Singh, who fails to recall that antibiotics are not indicated for acute bronchitis, and Randolph, who forgets to acquire milk on his way home. Although both are fully able to recall the crucial piece of information during the relevant period of time, neither does. Their ignorance does not excuse them from blame.

Another problem with Rosen's condition concerns cases of moral testimony (or deference). In such cases, an agent forms the true and justified belief that a certain act A is wrong by testimony from a reliable source. However, the source does not explain to the agent what makes A morally objectionable. As Alison Hills

<sup>5</sup> Why does acting from ignorance require ignorance of *every*, rather than *some*, wrong-making feature of one's act? Rosen (2008: 593) imagines an agent, Applebaum, who knows that his act puts a certain person, Botstein, at risk of harm, without knowing that it puts Botstein at risk of death. Although he is ignorant of a particular wrong-making feature of his action, Applebaum is still directly blameworthy for his action.

(2009) puts it, the agent may know that A is wrong, but she does not *understand* why (McGrath [2011] makes a similar point; see also Driver [2006], Hopkins [2007], and Jones [1999] on the phenomenon of moral testimony). The agent thus justifiably believes that A is wrong, while being ignorant about A's wrong-making features. Such an agent, it seems, would be blameworthy for performing A.

### 4.3 One Cannot Be Directly Culpable for One's Ignorance

Objection: Recall that the skeptical argument is based on the following premises:

- (1) If an agent S performs action A from ignorance, then S is culpable for the act only if S is culpable for the ignorance from which she acts.
- (2) An agent S is culpable for the ignorance from which she acts only if her ignorance is the upshot of some prior culpable act or omission.

On your view, an agent S may be *directly* culpable for her ignorance. You thus deny (2), but you have not discussed Rosen's and Zimmerman's arguments for (2).

In support of (2), Rosen writes that 'in the normal case, belief revision is a passive matter' (2004: 302). Furthermore, he adds, 'when I am passive with respect to an occurrence—when it merely happens in me or to me or around me—then I am responsible for the occurrence only if it is the (foreseeable) upshot of prior culpable activity on my part' (2004: 302). Similarly, Zimmerman (2008: 183–84) holds that one can be in direct control only of one's (physical or mental) actions. And belief, he adds, is not an action, and can thus at best be the *result* of an action. Hence, one can at best have *indirect* control of one's belief, perhaps 'by way of directly controlling a decision of which the belief is a consequence' (2008: 187).

Since one cannot be directly blameworthy for lacking a true belief that one's action is morally wrong, Rosen and Zimmerman conclude that one is culpably ignorant only when one culpably failed to do something that would have remedied one's ignorance. Culpable ignorance must be traceable to some past, original culpability.

Response: Before I address the most serious difficulty with Rosen's and Zimmerman's argument, two remarks are in order. First, Zimmerman's argument assumes that *direct* responsibility requires *direct* control. But this assumption is questionable. Suppose I fire a gun at Jones. On Zimmerman's view, I do not have direct control over the firing of the gun, because this action results from another action, namely, my pulling the trigger.<sup>6</sup> Now, my succeeding in firing the gun is not entirely up to me: the gun must function properly. However, assuming that I know that this additional condition of success *reliably* obtains, then I can be said to have control over my firing the gun. I have the capacity to fire the gun at will. And for this reason I can be held *directly* responsible for firing the gun it seems. Hence, although direct culpability requires control, it does not matter whether this control is direct

<sup>6</sup>As a matter of fact, on Zimmerman's view, I do not have direct control over my pulling the trigger either, for this action also results from another action, i.e., the mental action of deciding to pull the trigger (2008: 185).

or indirect. By contrast, a drunk driver is not *directly* blameworthy for hitting a pedestrian, not because he lacks *direct* control over his driving, but because he is unable to control the direction and speed of the car in a way that would enable him to reliably avoid hitting pedestrians. This suggests that the control condition on direct culpability should require not direct control, but merely the relevant capacity to reliably produce, directly or indirectly, the desirable outcome.

My second remark concerns Rosen and Zimmerman's assumption that responsibility requires *voluntary* control. We arguably have some form of control over our beliefs, even though that control is not voluntary. For example, although I incorrectly added two numbers, I could have figured out their sum. Because I have the ability to add correctly, it was, in a sense, 'up to me' to obtain the correct result. Perhaps this kind of control entails that one can be blameworthy for not forming a certain belief or for having a mistaken or unreasonable belief.

This may give rise to a different kind of objection against EC<sub>3</sub>. If one may be to blame for an unreasonable belief, then, it may be argued, the bank teller is *directly* culpable for forming a mistaken belief about the balance due the customer and is only *indirectly* to blame for reimbursing the customer with the wrong amount. In other words, this is a case of culpable ignorance: although she fulfilled her procedural epistemic obligations, the bank teller is to blame for her mistaken belief. Intuitively, she 'should have known better'.

This objection conflates two kinds of culpability, though. In forming her unreasonable belief about the balance due the customer, the bank teller violates an *epistemic* norm for which she might be epistemically blameworthy. But she clearly is not *morally* to blame for her epistemically unreasonable belief. Hence, this is not a case of morally culpable ignorance, and her moral culpability for reimbursing the customer with the wrong amount is not indirect.

Where does this leave us? I accept Rosen and Zimmerman's premise (2), according to which one cannot be *directly* (morally) culpable for one's ignorance. Culpable ignorance is *indirect* and due to a culpable violation of a procedural epistemic obligation. The more serious problem with their argument actually concerns premise (1): this premise is false, for one may culpably act from ignorance even though one is not (morally) culpable for the ignorance from which one acts. The bank teller does something morally wrong. Her wrongdoing is not due to an earlier moral wrongdoing. And she has no excuse for her wrongdoing since she is perfectly able to figure out the balance due and reimburse the customer accordingly. Similarly, Randolph is to blame for not getting the milk, even though he is ignorant of his wrongdoing. And his culpability is direct, for it is not traceable to an earlier moral culpability. Neither agent is culpably ignorant; yet, each is culpable for wrongly acting from ignorance. Thus, (1) is false.

## 5. Control: Further Objections and Responses

I will now address a pair of objections that concern the condition of control or the idea that blameworthiness requires control.

## 5.1 The Control Condition

Objection: On your view, an agent may be directly blameworthy for an action, even though she does not realize that her action has a certain wrong-making feature or will likely bring about a certain negative outcome. But how can the agent then be said to exercise control with respect to that feature or outcome? Since we do not exercise control over anything of which we are unaware, your view must relinquish the control condition. As Neil Levy puts it, ‘if moral responsibility requires control, then it requires that we know what we are doing’ (2005: 5).<sup>7</sup>

Response: I do not deny that responsibility requires control. But there are different kinds of control. In section 4.3, I remarked that we do have (nonvoluntary) control over our beliefs. This means that the bank teller does have control over what she is doing since, by assumption, she has the cognitive abilities necessary to form the correct belief about the balance due the customer. In this sense, even though she unknowingly did the wrong thing, it was under her control to do the right thing. She is to blame, because she failed to exercise that control successfully. Hence, contra Levy, although blameworthiness requires control, it does not require knowledge of wrongdoing.

## 5.2 Acting on a Reasonable Belief

Objection: Suppose an agent does have a reasonable *dispositional* belief that his doing A is morally permissible, but this belief plays no role in the production of A. Let us assume that we all have a moral duty to recycle glass (when possible). Diego’s friend tells him that in Greenville, the business in charge of recycling is so inefficient that recycling glass causes more harm to the environment than throwing it in the trash. Since his friend is usually reliable about such matters, Diego accepts his words. However, it turns out that in this case Diego’s friend is mistaken: the right thing to do in Greenville is to recycle. While in Greenville, Diego carelessly throws a glass bottle in the trash, without reflecting on the morality of what he is doing. Diego’s action is thoughtless. Intuitively, he is blameworthy for his wrongdoing. However, since he does have a reasonable (dispositional) belief that his action is morally acceptable, EC<sub>3</sub> would excuse him. This seems wrong.

Response: The point is well taken. To be excused from blame, the agent must act *based* on a reasonable belief that his action is morally permissible. Instead of EC<sub>3</sub>, we should have:

<sup>7</sup> Sher (2009) defends an account similar to mine. According to him, an agent S is blameworthy for an act A only if (1) S is aware of the wrongness of A or if (2) S’s failure to recognize the wrongness of A is both (a) substandard and (b) explained by S’s constitutive psychology. And whether a failure to recognize is substandard is a function of (i) S’s cognitive capacities and (ii) the moral requirements that apply to S. Like me, Sher invokes a counterfactual account of capacities à la Smith (2003). A proper treatment of Sher’s view is not possible here, but I should note two things. First, Sher’s conditions (1) and (2) are better combined into my condition EC<sub>3</sub>. The reasonableness of S’s failure to recognize the wrongness of her act is not sufficient for S to avoid culpability: S should have a reasonable belief that her act is permissible. Second, Sher contends (2009: 145) that his epistemic condition is incompatible with the thesis that no one is responsible for what is beyond one’s control. As I am about to explain, this contention is misguided.



- (EC<sub>4</sub>) Agent S is directly blameworthy for her wrongdoing A only if S's doing A is not based on a reasonable belief that her doing A is morally permissible.

Now, in typical everyday circumstances, the thought that what we are doing is morally permissible rarely occurs to us. This does not prevent us from acting blamelessly. Hence, the belief that her act is morally permissible need not be present to the mind of the blameless agent. However, her act must be *regulated* by that belief. This means that had the agent not believed (or had she stopped believing) that the act is morally permissible, she would not have performed the act. Diego is culpable for throwing a glass bottle in the trash, because his belief that doing so is morally acceptable does not in any way constrain his action: he would have performed it even if he had not had the belief. In EC<sub>4</sub>, the locution 'based on a belief' should thus be understood, broadly, to include actions that are regulated by that belief.

## 6. Conclusion

Morality imposes what I have called *primary* norms. A morally wrong action is one that violates such a norm. But agents (as opposed to actions) are also the targets of our moral assessments: they may be deemed blameworthy (or praiseworthy) for their actions. A judgment of culpability, I have argued, is based on a *secondary* norm requiring agents to do their reasonable best to respect primary moral norms. Agents are thus to blame in cases when their failure to respect morality is an underperformance, that is, when it is a failure to exercise their capacities successfully.

This account of moral responsibility has several virtues. First, it blocks a powerful skeptical argument. Crucially, it exposes the argument's mistaken premise according to which culpable wrongdoing has to be deliberate. Second, the account consists in a very simple norm: to be responsible is to do what one can to act morally. Third, I have shown, the account respects our everyday judgments of moral as well as prudential culpability. Finally, the account offers a unified treatment of various forms of moral culpability. Blameworthy actions and omissions include deliberate wrongdoings, wrongful omissions due to apathy, wrong actions resulting from cognitive lapses and blunders. These actions or omissions are in many respects dissimilar. However, they are all culpable because they are all instances of underperformances, that is, of failures to respect morality to the extent of one's abilities. I have distinguished among three types of underperformance.

First, an agent may fail to deploy his cognitive abilities adequately to figure out what the right thing to do is. This happens when his attempt to act (or omission) is not based on a reasonable belief that the act (omission) is morally permissible. Most of the culpable agents we have considered fall into this category. The bank teller who miscalculates the balance due a customer and Dr. Singh, who fails to recall the crucial information learned at the conference, are both blameworthy because they

lack an epistemically reasonable belief that they are doing something permissible. Diego, the careless nonrecycler, does reasonably believe that putting glass in the trash is acceptable, but since his putting glass in the trash is not based on (in this case, regulated by) that belief, he is to blame for his wrongdoing. Similarly, we may suppose that the police officer who bludgeons the demonstrator dispositionally knows (or at least reasonably believes) that his act is wrong. And like Diego, his action is not properly regulated by that belief. Randolph, who fails to acquire milk, is in some respects like Diego. He clearly had previously planned to get milk, but during the relevant period of time, he failed to attempt any action based on his dispositional belief that he ought to stop at the store. Alternatively, Randolph can be seen as lacking a reasonable belief that there is nothing wrong with his actions during that portion of time. On both readings, he fails to perform according to his cognitive capacities.

Second, an agent may not exploit his volitional capacities adequately. A person who fails to meet her moral obligation because she does not appropriately tap into her motivational capacities is to blame for her inaction. Similarly, a parent is to blame for snapping at his child if his action results from a failure to exercise his self-control successfully.

Third, an agent may not successfully use his motor skills to do the right thing. An experienced surgeon who bungles a routine operation has a reasonable belief about what she ought to do, attempts to act on the basis of that belief, but fails to act according to her physical abilities. Note, however, that this proposed division into three categories is a little contrived, since most of our actions require the simultaneous deployment of cognitive, volitional, and motor abilities.

MARTIN MONTMINY  
UNIVERSITY OF OKLAHOMA  
[montminy@ou.edu](mailto:montminy@ou.edu)

## References

- Austin, J. L. (1970) *Philosophical Papers*. Oxford: Oxford University Press.
- Block, Ned. (1981) 'Psychologism and Behaviorism'. *Philosophical Review*, 90, 5–43.
- Clarke, Randolph. (2014) *Omissions: Agency, Metaphysics, and Responsibility*. Oxford: Oxford University Press.
- Driver, Julia. (2006) 'Autonomy and the Asymmetry Problem for Moral Expertise'. *Philosophical Studies*, 128, 619–44.
- Fara, Michael. (2008) 'Masked Abilities and Compatibilism'. *Mind*, 117, 843–65.
- Fischer, John Martin, and Mark Ravizza. (1998) *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge, UK: Cambridge University Press.
- FitzPatrick, William. (2008) 'Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge'. *Ethics*, 118, 589–613.
- Ginet, Carl. (2000) 'The Epistemic Requirements for Moral Responsibility'. *Philosophical Perspectives*, 14, 267–77.
- Haji, Ishtiyaque. (1997) 'An Epistemic Dimension of Blameworthiness'. *Philosophy and Phenomenological Research*, 57, 523–44.
- Harman, Elizabeth. (2011) 'Does Moral Ignorance Exculpate?' *Ratio*, 24, 443–68.

- Hills, Alison. (2009) 'Moral Testimony and Moral Epistemology'. *Ethics*, 120, 94–127.
- Hopkins, Robert. (2007) 'What Is Wrong with Moral Testimony?' *Philosophy and Phenomenological Research*, 74, 611–34.
- Jackson, Frank. (1986) 'A Probabilistic Approach to Moral Responsibility'. In R. Barcan Marcus, G. J. W. Dorn, and P. Weingartner (eds.), *Logic, Methodology and Philosophy of Science VII* (Amsterdam: Elsevier Science Publishers), 351–65.
- Jones, Karen. (1999) 'Second-Hand Moral Knowledge'. *Journal of Philosophy*, 96, 55–78.
- Levy, Neil. (2005) 'The Good, the Bad and the Blameworthy'. *Journal of Ethics and Social Philosophy*, 1, 1–16.
- Levy, Neil. (2011) *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. Oxford: Oxford University Press.
- McGrath, Sarah. (2011) 'Skepticism about Moral Expertise as a Puzzle for Moral Realism'. *Journal of Philosophy*, 108, 111–37.
- McKenna, Michael. (2012) *Conversation and Responsibility*. Oxford: Oxford University Press.
- Mele, Alfred. (2003) 'Agents' Abilities'. *Noûs*, 37, 447–70.
- Robichaud, Philip. (2014) 'On Culpable Ignorance and Akrasia'. *Ethics*, 125, 137–51.
- Rosen, Gideon. (2004) 'Skepticism about Moral Responsibility'. *Philosophical Perspectives*, 18, 295–313.
- Rosen, Gideon. (2008) 'Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility'. *Journal of Philosophy*, 105, 591–610.
- Sher, George. (2009) *Who Knew? Responsibility without Awareness*. Oxford: Oxford University Press.
- Smith, Angela. (2005) 'Responsibility for Attitudes: Activity and Passivity in Mental Life'. *Ethics*, 115, 236–71.
- Smith, Holly. (1983) 'Culpable Ignorance'. *The Philosophical Review*, 92, 543–71.
- Smith, Michael. (2003) 'Rational Capacities, or: How to Distinguish Recklessness, Weakness, and Compulsion'. In S. Stroud and C. Tappolet (eds.), *Practical Irrationality* (Oxford: Clarendon Press), 17–38.
- Vihvelin, Kadri. (2004) 'Free Will Demystified: A Dispositional Account'. *Philosophical Topics*, 32, 427–50.
- Whittle, Ann. (2010) 'Dispositional Abilities'. *Philosophers' Imprint*, 10, 1–23.
- Zimmerman, Michael. (1997) 'Moral Responsibility and Ignorance'. *Ethics*, 107, 410–26.
- Zimmerman, Michael. (2008) *Living with Uncertainty*. Cambridge, UK: Cambridge University Press.