

# Explaining away epistemic skepticism about culpability

Gunnar Björnsson, Stockholm University

## 1. Epistemic skepticism about culpability

Ignorance often excuses. Consider:

*Vera and the Vase:* Vera opens one of the front doors of the small museum to get to her car. As she opens the door, it pushes a tall vase off balance. The vase, temporarily placed just outside that door by a cleaner, falls over and shatters as it hits the floor, rendered irreparable.

Here are two ways of filling in the details of the case:

*Suspecting Vera:* Walking towards the door, Vera was aware that a vase might be standing just outside—she saw a reflection of the vase in an open window next to the door—and knew that she should peek out to make sure or leave through the other door. Not wanting to take the extra time, however, she hoped for the best and opened the door.

*Unsuspecting Vera:* Vera had no reason to suspect that anything bad would happen if she opened the door, and so no reason to be extraordinarily careful or to leave through another door.

Suspecting Vera, it seems, can be blamed for taking the risk and destroying the vase. Not so for Unsuspecting Vera: she took her behavior to be unproblematic, and was justified in doing so. It could still make sense for her to feel bad about the outcome, and for the museum personnel to be upset. But the relevant emotional reactions are not those characteristic of moral blameworthiness—guilt and indignation—but instead something like what Bernard Williams (1976) calls “agent regret” and its second or third person counterparts. Similarly, Vera might have some (*pro tanto* or *prima facie*) duty to help amend the situation, if nothing else by notifying the museum personnel, but such duties can arise for outcomes for which one is not morally responsible or morally to blame.

Based on examples such as these, it is natural to conclude that when one acts based on the (implicit or explicit) belief that one’s action is justifiable and one cannot be blamed for that belief, one also cannot be blamed for one’s action. By the same token, it also seems plausible that if a belief has been formed based on the (implicit or explicit) assumption that the way in which it was formed was justifiable and one cannot be blamed for that assumption, one cannot be blamed for the belief. If so, it might seem that any blameworthiness will have to be grounded in acts (including acts of belief formation) going against one’s belief that what one is doing is unjustifiable—acts of *akrasia*. For illustration, consider another variation of *Vera’s Vase*:

*Wrongheaded Vera:* Vera was aware that a vase might be standing outside the door but thought that she was entitled to risk damaging it. If the museum were placing fragile objects in such a way, it shouldn’t be *her* problem.

At first glance, Wrongheaded Vera seems blameworthy for her action: she acts on a self-serving normative understanding of the situation, giving too little weight to the possible damage she might cause. But given the requirement that blameworthiness be grounded in acts of *akrasia*, and given that she took her action to be justifiable, her blameworthiness depends on whether she had previously taken what she then believed was an unjustifiable risk of reaching an erroneous moral conclusion. Moreover, akratic acts seem relatively rare, and akratic acts of unjustifiably reaching mistaken moral conclusions even more so. If nothing else, it can be very hard to know whether others' wrongful behavior is ultimately the upshot of *akrasia*. Perhaps, then, we should be highly skeptical about most attributions of blame—or so a number of people have recently argued (see e.g. Zimmerman 1997; Rosen 2003; 2004; 2008; Zimmerman 2008; Levy 2009; 2011, ch. 5; for criticism, see e.g. FitzPatrick 2008; Harman 2011).

While reflection on cases like *Suspecting Vera* and *Unsuspecting Vera* might seem to support this skeptical conclusion, other cases of unwitting wrongdoing suggest that blameworthiness is much more common. Consider:

*Scatterbrained Vera*: Vera had seen the vase reflected in the open window next to the door and briefly noted how beautiful it was, but failed to integrate the information about its location into the spatial map guiding her action.

Having failed to put two and two together, as it were, it might make sense for Vera not only to feel agent regret, but also to blame herself for what happened: “I should have realized”. The following paraphrase of an example from Randolph Clarke illustrates the same point:

*Forgotten Milk*: You had promised your spouse to buy milk on the way home before the stores closed, but your thoughts drifted and you forgot. Your failure to remember was neither a case of *akrasia* nor the result of some blameworthy act, such as a failure to set the alarm on your phone, for you almost never forget things like these. Still, you can rightly be blamed for the lack of milk for breakfast (Clarke 2014).

In light of this and similar cases, Clarke suggests that an agent can be responsible for ignorant wrongdoing when her lack of awareness “falls below a cognitive standard that applies to her given her cognitive and volitional abilities and the situation she is in” (Clarke 2014: 303).<sup>1</sup>

I agree with Clarke that blame can be appropriate in such cases, and exactly because the agent has fallen short of a cognitive standard. But contrary to what Clarke seems to assume, the appropriate sort of blame is not obviously *moral* blame. Rather, it seems to be a non-moral form of what we might call “skill blame”. Noticing and locating objects based on visual information, remembering things, or solving mathematical problems are activities that can be performed more or less skillfully, with greater or lesser excellence, and in all these areas we can be blamed or credited for resulting successes and failures. However, while you might kick yourself for failing to notice the location of an object or remember something, or for failing to

---

<sup>1</sup> Similar cases are discussed at some length by George Sher (2009). In a proposal that is structurally similar to Clarke's, FitzPatrick (2008) suggests that blameworthiness can be grounded in the agent's display of epistemic vices.

solve a mathematical puzzle, characteristically *moral* modes of blame involving guilt and indignation seem inappropriate. Moreover, they seem inappropriate even if the predictable outcome of your failure was horrific and morally highly significant:

*Fatal Failure:* To disarm the ticking bomb in a densely populated building, you had to quickly solve a complex mathematical puzzle. Unfortunately, in spite of a highly focused effort, you failed to prevent the lethal outcome.

Suppose that you had the capacity so solve the puzzle, and that your failure fell short of a cognitive standard that applied to you in the situation: you simply made a stupid mistake. Then agent regret for your involvement in the outcome and skill blame for your failure to solve the puzzle might both be appropriate, but you do not deserve *moral* blame for the tragic outcome. It wasn't due to any moral shortcoming of yours. But similarly, if your failure to keep the promise to your spouse wasn't due to any moral shortcoming, it is unclear why you would be morally to blame for it.<sup>2</sup>

I have suggested that the requirement that blameworthiness be grounded in some act of *akrasia* is untouched by the blame appropriate in cases like *Forgotten Milk*. I also believe that forms of epistemically based skepticism are deeply intuitively compelling. Moreover, these forms of skepticism seem independent of standard sources of responsibility skepticism, such as worries that putative objects of responsibility are the upshot of luck, or determined by factors outside of our control.

In this paper, however, I argue that radical epistemic responsibility skepticism is mistaken. In section 2, I sketch what I take to be the most promising general account of moral responsibility, blame and credit, building on the broadly Strawsonian idea that agents are morally blameworthy for objects that relate in the right way to a *substandard quality of will*, and the idea that responsibility is a certain kind of *explanatory* relation. In section 3, I explain how this account deals with central cases seemingly supporting the *akrasia* requirement, and how it avoids some objections to quality of will accounts recently raised by Gideon Rosen. But the intuitive worry brought out by these objections go beyond what is directly handled by a plausible quality of will account. In section 4, I spell out this remaining worry, argue that it shares a crucial structure with familiar skeptical or incompatibilist worries, and suggest that strategies for explaining away the intuitive pull of the latter are applicable in this case too.

---

<sup>2</sup> On quality of will accounts of blameworthiness such as that outlined below, you might be blameworthy for your failure to keep your promise if it is due to a failure to care appropriately about relevant matters. But Clarke's example is designed to suggest that your shortcoming is merely cognitive.

My claim that you do not deserve *moral* blame in *Forgotten Milk* is partly stipulative. It might be appropriate for you to feel some (mild) agent regret for not keeping your promise (and for your spouse to be correspondingly (mildly) annoyed with you), and you might deserve non-moral skill blame for your memory lapse. You might also have a special obligation to mend the situation and to explain yourself. These or very similar features would be present if the forgetfulness had been due to a moral shortcoming on your part, and they are present because something morally bad happened: you did not keep your promise. In light of this, it is natural to say that you are morally to blame for the failure. However, such features can also be present in cases where you are clearly not morally culpable, such as that of *Unsuspecting Vera*.

## 2. The *explanatory quality of will* account of moral responsibility

This section outlines what I take to be the most promising account of moral responsibility, blame, and credit—the *Explanatory Quality of Will* account (EQW)—and indicates how it yields unified explanations of a variety of familiar facts about moral responsibility. The presentation falls far short of a full defense, but I have motivated aspects of the account in detail elsewhere,<sup>3</sup> and elements that the account has in common with other quality of will accounts have received thorough defenses in the literature.<sup>4</sup>

The EQW builds on two main elements. The first is a general account of the sort of agent-directed responsibility, blame, and credit that grounds criticism and praise in both moral and non-moral domains. *Crediting* or *blaming* someone, in the senses that concern us here, are modes of positive or negative evaluation that involve seeing the agent as responsible for some good or bad object, the object *for which* the agent is responsible and credited or blamed. The basic idea is that such evaluations are fitting if and only if the good or bad object of responsibility is the upshot of a corresponding shortcoming or excellence of the agent. A *shortcoming* here is some substandard aspect of the agent—some aspect falling short of what can be properly demanded of the agent—and an *excellence* is a corresponding superstandard aspect. For example, the student who arrives at the wrong answer to a mathematical question deserves math-related blame for this insofar as she made the error because she fell short of what could be properly demanded of her in terms of mathematical understanding, say, and the student correctly answering the question deserves math-related credit if solving it required her to surpass such demands. Similarly, a psychologist deserves professional blame for misdiagnosing a patient insofar as she did so because her skills or mind-set fell short of what can be properly demanded of her in her role, and a basketball player deserves basketball-related credit for blocking a shot insofar as succeeding required surpassing what could be properly demanded of her in that role.

The *amount* of blame or credit given to someone typically depends at least partly on the *value* of the object of blame or credit. The player who makes the game-winning goal from a difficult angle gets more credit than the player who delivers a similarly difficult but less consequential pass, and a bumbling detective who sinks an investigation by destroying crucial evidence gets more blame than one whose equal amount of incompetence destroys evidence

---

<sup>3</sup> For an overview and more detailed application than can be offered here, see Björnsson Forthcoming-a; for some of my earlier discussion of this kind of account, see Björnsson 2011; Björnsson and Persson 2012; 2013; Björnsson 2014a; b; Björnsson and Pereboom 2016. The current proposal is in many ways similar to Nomy Arpaly and Timothy Schroeder’s recent (2014) account of blame- and praiseworthiness of actions specifically, but covers blame and credit for omissions and non-agential events as well as for actions, and understands quality of will in relation to what can be properly demanded of agents in line with the general account of blame and credit.

<sup>4</sup> For the broadly Strawsonian idea that moral blame and credit depends on the agent’s concern with what is morally important, see especially Nomy Arpaly’s work (Arpaly 2003; 2006; Arpaly and Schroeder 2014; also see McKenna 2012). The idea that responsibility for an action is primarily a matter of how that action came about is familiar from John Fischer and Mark Ravizza (1998) and others working in the tradition of Harry Frankfurt’s (1969) “Alternate Possibilities and Moral Responsibility”.

of less importance. (The value of an action might itself depend on the shortcoming or excellence involved in or displayed by that action. In sports, a move requiring high degrees of skill might itself be seen as valuable and as an object of credit even if it does not in the end contribute to victory, and something is a bad act of recklessness only insofar as the act displayed a substandard lack of concern with certain risks.)<sup>5</sup>

Amounts of blame and credit also depend on how much of a shortcoming or an excellence was required to explain the outcome. The player who scored in circumstances requiring a high degree of excellence gets more credit for doing so than the player scoring in more favorable circumstances, and the student failing to answer the mathematical question gets more blame if the question was easier and failure required more of a shortcoming in mathematical ability. Notice that what matters here is not the agent's degree of excellence or shortcoming, but only the degree required to explain the outcome. If the same degree of mathematical ability can be required of two students who have solved a certain mathematical problem, they get the same amount of credit for solving the problem even if one of them has a higher ability.

What can be properly demanded, and thus what counts as a shortcoming or excellence, presumably varies with context and with the capacity of the agent. Lower degrees of mathematical ability can be properly demanded of a novice or someone operating under a great deal of stress. In such cases, failure to correctly answer a given mathematical question might involve less of a shortcoming or none at all, and so yield lower levels of blame or none at all, and success in correctly answering the question might require a higher ability compared to what can be properly demanded, and so more credit, or credit where none would otherwise be due.

Shortcomings and excellences that ground blame and credit are negatively and positively valued at least in part because of their tendency to bring about correspondingly negatively and positively valued objects. Sport skills are valued because they contribute towards winning, mathematical abilities because they contribute towards mathematical understanding and problem solving, the skills of a detective because they contribute to the identification of culprits and gathering of evidence that holds up in court, a cool mind because it contributes to sensible action under pressure, wit because it contributes to relief and enjoyment, and so forth. (Low levels of skill or a skittish or dull mind are correspondingly negatively valued because they have the opposite tendency.) Importantly, such tendencies to bring about good or bad are relative to circumstances and involve characteristic ways of playing out. A "normal" way for an excellence to lead to something good or a shortcoming to lead to something bad is whatever way it typically leads to that sort of thing under the relevant circumstances.

This sort of *normality* matters for credit and blame. Suppose that a criminal seeking to create trouble for a competitor makes sure to tip off the best detective in the police force.

---

<sup>5</sup> Some will insist that the only thing that has positive or negative *moral* value is the expression of a morally super- or substandard state of mind, and that the amount of blame depends solely on that state and not on resulting outcomes. Though I think that this view is mistaken, often based on ultimately unfounded worries about outcome luck (see Björnsson and Persson 2012), it is strictly speaking compatible with the general account of blame and degrees of blame proposed here.

Suppose further that, as a result, that detective identifies the competitor as the perpetrator and collects the easily available evidence needed for a conviction. Here the detective’s high level of ability was a crucial part of why she solved the case—she would not have been tipped off without it. But her level of excellence did not explain the result in the normal way that it explains such outcomes, as she only had to use her basic skills. Consequently, she does not deserve credit for the outcome corresponding to the degree of excellence that was necessary for it. Similarly for blame: Suppose that the colleague of a bad detective actively sabotages an investigation where the bad detective is making good progress because the colleague thinks that such a bad detective doesn’t deserve success. Then the shortcomings of the bad detective were part of why the investigation failed, but not in the right way to attach the blame to her.<sup>6</sup>

Notably, tendencies for agential aspects to bring about good or bad can be *general* or *more specific*. A student might have a high general ability to solve a certain kind of mathematical problem, but suddenly find herself unable to do so. (We can think of strengths of tendencies as corresponding to proportions of circumstances within a certain set in which the agential aspect has the relevant outcome for normal reasons. The student might successfully execute in most of a broad set of circumstances, but have a low or zero rate of success within some small subset.) Correspondingly, demands on agential aspects might concern aspects with both general and specific tendencies. An agent might thus live up to the general demands but still fall short of demands pertaining to more specific circumstances.<sup>7</sup> As a result, an overall good person might be to blame for failures that are “out of character”.

The proposal outlined here can be summarized as follows:

EXPLANATORY BLAME (CREDIT): X deserves blame (credit) for Y if and only if Y is bad (good) and explained in a normal way by X’s corresponding aspect falling below (above) what could be properly demanded of X.

DEGREE FROM DEVIATION: The degree of blame (credit) X deserves for Y depends on the value of the outcome and on how much of a deviation of the agential aspect from what can be properly demanded is required in the normal explanation of the object of blame (credit).

---

<sup>6</sup> In *Fatal Failure*, your mathematical shortcoming explained your failure to solve the puzzle in a normal way. But it might seem that it also explained the lethal outcome in a normal way: a failure of skill leading to a bad outcome that an application of the skill was meant to solve. Here it is important to remember that standards are set in relation to practices concerned with specific values. While you did fall short with respect to standards of mathematical skill relating to *mathematical* understanding and problem solving, *Fatal Failure* does not suggest that you fell short with respect to standards that are grounded in a general interest to achieve various practical outcomes or any specific interest to avoid disasters. (For this reason, the degree or math-related blame deserved is affected by the mathematical badness of not solving the puzzle, not by the badness of the fatal outcome. See DEGREE FROM DEVIATION.)

<sup>7</sup> Demands for aspects that tend to do good under specific circumstances might be in place exactly because it is natural to fall short in such circumstances even if one tends to do good generally, and procedures might be in place to help agents satisfy such demands, as when personnel practice dealing with emergency situations.

Above I said that blaming and crediting involves taking the subject to be responsible for the object of blame or credit. Given what has been said, the relevant notion of responsibility would thus be:

EXPLANATORY RESPONSIBILITY: For X to be responsible for Y is for Y to fall in an evaluative dimension and be explained in normal ways by a corresponding aspect of X.

Notice that on this account, responsibility (and blame or credit) does not require that the agent is responsible for the agential aspects that ground her responsibility.

These explanation-related accounts of responsibility and degrees of blame and credit provide the first part of the Explanatory Quality of Will account. The second part follows a tradition from P. F. Strawson's (1962) "Freedom and Resentment" in identifying the agential aspect relevant for *moral* credit and blame specifically with the agent's quality of will: this is the sort of agential aspect that grounds the sort of crediting and blaming that involves attitudes of moral gratitude, on the one hand, and guilt and indignation, on the other.<sup>8</sup> Together with EXPLANATORY BLAME (CREDIT), the quality of will account yields:

MORAL EXPLANATORY BLAME (CREDIT): X deserves moral blame (credit) for Y if and only if Y is morally bad (good) and explained in a normal way by X's quality of will falling below (above) what could be properly morally demanded of X.<sup>9</sup>

I take an agent's quality of will to be a matter of how well she *cares* about what is morally important, where caring about something in the relevant sense involves being disposed to pay particular attention to information relevant for promoting or not obstructing the object of caring and to have one's behavior be guided by such information. In the case of *Fatal Failure*, where you failed to prevent a deadly explosion because you made a mathematical mistake, you might deserve math-related blame if your mathematical savvy fell short of what could be properly demanded of you, but you would only be morally to blame for the outcome if it was the upshot of your failure to care appropriately. (In what follows, when I talk about blame and credit, I will mean *moral* blame and credit unless otherwise stated.)

In principle, the reliance on a notion of what degrees of caring *can be properly demanded* leaves room for a variety of conditions on responsibility discussed in the literature. One could

---

<sup>8</sup> In focusing on these agential aspects, I am not denying the importance of other aspects, and other evaluative reactions (see Shoemaker 2013).

<sup>9</sup> This falls short of a full account of deserved blame and credit in three ways: First, it says nothing about what counts as morally good or bad, thus qualifying as object of credit or blame. Candidates include, among others, wellbeing, harm, certain kinds of actions, certain relationships, and certain character traits or states of mind, all of which might be understood as agent-relative or agent-neutral values. Second, it says nothing about what it can be demanded that a given agent cares about, or to what extent. This is presumably closely connected to what counts as morally good or bad, but will also depend on such things as the agent's capacity for caring about these things. Finally, I take there to be more than one kind of moral credit. One is captured by MORAL EXPLANATORY BLAME (CREDIT) and DEGREE FROM DEVIATION, but another, equally important, instead requires that the agent put some significant effort or sacrifice into doing moral good, even if that sacrifice was itself required rather than supererogatory (for examples, see Massoud 2016). I set this issue to the side for reasons of space.

say, for example, that for a certain degree of caring to be properly demanded of an agent at a time, the agent must have had a certain kind of incompatibilist control over whether she would care to that degree at that time. I bracket such requirements, however, as those who take issues of ignorance to motivate responsibility skepticism typically take this motivation to be separate from what drives traditional incompatibilist skepticism or associated demands of ultimate control. What any account of responsibility, blame, and credit should be sensitive to, however, is the agent's ability to grasp the relevant concepts and moral reasons, and perhaps also her ability to be motivated by morally relevant features. Moreover, since such abilities come in degrees, it is hard to avoid the conclusion that the extent to which it can be morally demanded that someone care about certain features also comes in degrees. In sum, I will thus assume that:

DEMAND FROM CAPACITY: The degree to which it can be properly demanded that X cares about something depends only on X's capacity for so caring, not on how X came to care to the degree she does.

I will understand X's capacity for caring about something to a certain degree as a matter of how easy it would be for X to do so. I also take caring to some degree to be less easy the more X would have to be committed to so caring in order to achieve and uphold it given X's relevant cognitive and motivational capacities and other relevant circumstances.

Together, EXPLANATORY MORAL BLAME (CREDIT) and DEGREE FROM DEVIATION straightforwardly account for central aspects of moral blame and credit (for more elaborate discussion, see Björnsson Forthcoming-a). Consider first the account of how partial but insufficient justification of a wrongful act lowers the degree of blame. Here, the explanation of a bad act (a lie, say, or an act of violence) will require less of a deviation in quality of will when there is *some* (though insufficient) justification than when all justification is absent. Second, consider how it accounts for the role of force, physical constraints, and threats as partial or full excuses or justifications. In typical cases of completely overwhelming physical force, there is no normal role for the agent's quality of will in explaining the outcome. If resisting the physical force is possible, the resulting behavior might to some extent be explained in normal ways by the agent's quality of will, but the higher the difficulty and cost, the less of a deviation in quality of will might be required to explain why the agent performed the action. Likewise for threats: they can provide the agent with strong valid reasons to perform the action and decrease the extent to which the explanation of the action requires that the agent's quality of will fall short of what can properly be demanded of her. In the case of severe threats, even someone satisfying those demands would perform the action, rendering any ill will explanatorily superfluous and the agent blameless for that action.<sup>10</sup>

Finally, consider the account of why children might deserve less blame and more credit for a given kind of action than adults, and why someone who is blamelessly inebriated or exhausted might deserve less blame and more credit for a given act than someone in a normal psychological state. Given that children will often have less developed cognitive and emotional

---

<sup>10</sup> If performing the action *with a bad motive* is itself morally bad, the agent might still be to blame for that, even if she cannot be to blame for performing the action.



capacities than adults, the levels to which we can properly demand that they care about morally important matters are lower, meaning that certain good actions that are par for the course for an adult requires superstandard quality of will for a child. Similarly, the levels of general caring about morally important matters that we can properly demand of people might not be enough to guarantee that they behave well when inebriated or exhausted, and the levels of caring that can be properly demanded *specifically* for such situations might be lower exactly because inebriation lowers cognitive and emotional capacities.

### 3. The *Explanatory Quality of Will* account and ignorant wrongdoing

Consider next what EQW says about the excusing role of ignorance, beginning with variations of *Vera and the Vase*. *Unsuspecting Vera* had no reason to worry that opening the door would topple a vase, and thus seemed blameless for the outcome. This is easily explained given EXPLANATORY MORAL BLAME (CREDIT): the outcome is not explained by any deviation of her quality of will from what can properly be demanded of her. *Suspecting Vera*, by contrast, realized that opening the door posed an unjustifiable risk, but went ahead, hoping for the best. According to EXPLANATORY MORAL BLAME (CREDIT), this was blameworthy given the plausible assumption that her decision stemmed from a failure to care sufficiently about avoiding risks to others. *Wrongheaded Vera* thought that she was entitled to risk the museum property because they had put their property in danger's way themselves. Again, EXPLANATORY MORAL BLAME (CREDIT) offers an intuitively appealing account of why she is to blame for destroying the vase: she willingly risked damaging museum property because she cared too little about the consequences of her own actions—too little compared to what could properly be demanded of her. The fact that she thought that she was justified in taking the risk does not in itself show that she cared about morally relevant matters as can be properly demanded of her. *Scatterbrained Vera*, finally, might be to blame for not realizing that there might be a vase just outside the door, but since this failure was plausibly due to a purely cognitive shortcoming and not due to substandard quality of will, she would not be *morally* to blame for the outcome.

In this regard, *Scatterbrained Vera* is different from other cases where agents fail to notice possible effects of what they are doing (Björnsson Forthcoming-a):

*Knockout:* Leaving the room, Victor pushes the door open quickly and with great force, inadvertently knocking unconscious the person just about to open the door from the busy corridor outside. At the moment of action, it didn't cross Victor's mind that opening the door in that way might hurt someone, though he would have realized this if the question had come up.

*Badgering:* Badgering his interlocutor with questions, Victor fails to pick up readily available signs of severe distress in her face and to adjust his behavior. His insensitive interrogation, happening at an unfortunate time, causes her great discomfort.

Victor never considered the possibility that he might knock someone unconscious by opening the door, or cause great discomfort by his insistent questions. Still, most would be inclined to

assign at least *some* blame to Victor for the effects (though less than if he had brought these effects about knowingly). It is clear how EXPLANATORY MORAL BLAME (CREDIT) leaves room for these judgments: in both cases, the failure to access the relevant information, from memory and the visual field, respectively, might plausibly be explained by a failure to care as can be reasonably demanded of Victor, about risks he is imposing on others, and about their wellbeing. It similarly accounts for moral blameworthiness for forgotten promises when the forgetting is plausibly explained by a substandard quality of will (as in versions of *Forgotten Milk* where your forgetfulness is the upshot of a failure to care appropriately about your spouse, or about promise keeping). Finally, it is clear how DEGREE FROM DEVIATION accounts for the sense that Victor is considerably less blameworthy than he would have been if he had willfully ignored the risk of forcefully pushing a door open into a busy corridor, or willfully caused his interlocutor great discomfort. Doing such things requires more of a deviation in quality of will from what can be properly demanded.

In broad outline, these applications of a quality of will account of moral blameworthiness are familiar from the literature,<sup>11</sup> and it is clear how they help the account avoid one kind of epistemically based skepticism about culpability, by foregoing the requirement that blame must be grounded in acts of akrasia. This is not to say that it is easy to determine whether someone deserves moral blame given EQW. What sorts of quality of will can properly be demanded of people is itself a difficult issue, and humans have a well-known tendency to over-attribute actions to character traits rather than circumstantial factors. Moreover, whether an action or omission is explained by a certain quality of will of an agent typically depends on the agent's beliefs and accessible information, both of which might be hard to ascertain, at least for an outsider. But compared to cases of akrasia, actions or omissions explained by substandard quality of will seem much more common, and less difficult to identify.

In foregoing the akrasia requirement on blameworthiness, EQW's obviously runs contrary to arguments in support of that requirement. For example, Michael Zimmerman (1997; 2008, ch. 4) assumes that acts from moral ignorance are non-culpable if the ignorance is non-culpable and argues that one is not in direct control of one's moral ignorance, and so cannot be culpable for it unless it is the upshot of some wrongdoing over which one was in direct control. This wrongdoing would thus have to be an act of clear-eyed akrasia. Levy (2011, ch. 5) argues (a) that, except in cases of clear-eyed akrasia, agents who do wrong do so in accordance with their internal reasons, and (b) that it cannot be demanded of agents that they act in ways incongruent with their internal reasons.<sup>12</sup> Against these arguments, the independent support available for EQW and the fact that it does a better job accounting for blameworthiness in cases like *Knockout* and *Badgering* provide independent reason to reject premises in these arguments. In response to Zimmerman, this gives us reason to say that agents in cases like *Knockout* and *Badgering* have the relevant kind of control over their ignorance in cases where it is the normal upshot of substandard caring. In response to Levy, it gives us reason to say that the moral demands that ground blame for non-akratic action are

<sup>11</sup> See e.g. Arpaly 2003: 160–62; A. M. Smith 2005; H. M. Smith 2011.

<sup>12</sup> Gideon Rosen (2003: 79–80; 2004: 306) sketches an argument very similar to Levy's. For critical discussion of Levy's version, see Robichaud 2014.

not demands that agents act contrary to their internal reasons, but that they care sufficiently about morally important matters (and so have different internal reasons).

It might be thought, though, that quality of will accounts like EQW run into problems with a family of cases provided by Gideon Rosen. In these cases, agents do bad things out of a lack of proper regard for others, but blameworthiness seems at least significantly mitigated by difficult epistemic circumstances that cannot themselves be blamed on the agents.<sup>13</sup> Among the characters that Rosen asks us to consider is the Hittite lord who “buys and sells human beings, forces labour without compensation, and separates families to suit his purposes (Rosen 2003: 64)”. The lord knows fully well that his actions degrades and harms the slaves and has no relevant false non-moral beliefs about the nature of the slaves that justifies the practice, but in his ancient Near East society no one has yet thought to object to slavery on moral grounds. Another character is the 1950s US father who supports the college education of his sons but not his daughters, unreflectively taking this to be obviously right because he was raised to do so, and because people he takes seriously take it to be obviously right: “an undefended axiom of moral common sense” (ibid.: 66–69; cf. Calhoun 1989). As Rosen points out, it would have taken an extraordinary feat of moral invention for these characters to see the wrongness of their ways given their cultural surroundings. For this reason, serious moral blame might seem out of place. The victims can rightly be upset with the actions and there is room for serious moral criticism, but when one takes into account the circumstances in which these agents found themselves for no fault of their own, it might seem deeply problematic to say that they deserve our indignation.<sup>14</sup> And yet there is no doubt that these agents are acting from lack of proper regard for some of their fellow human beings.

I share Rosen’s intuitive sense that moral blame is significantly undermined in cases like these. But at least at first glance, there seem to be reasons to think that this is compatible with EQW.

The most obvious is that in order to behave like these characters do *in a modern western social setting*, one must not only care insufficiently about equal opportunity or the lives of the slaves, but do so in spite of being repeatedly prodded to reconsider and reevaluate. This, it seems, would require a considerably worse quality of will than what would be required in a context without social pressure to change one’s attitudes and behaviors, involving more of a determination not to take certain considerations into account. Though this falls short of showing that the Hittite lord or 1950s sexist father would be blameless given EQW, it does suggest that they would deserve less blame than their contemporary counterparts. Relatedly, as Nomy Arpaly and Timothy Schroeder (2014: 184) point out, others’ favorable attitudes towards a certain kind of action can cloud one’s understanding of what one is doing when performing it, making it harder for the Hittite lord to see clearly how his actions relate to such

---

<sup>13</sup> Rosen 2003: 72–73; 2008: 607–09. Cf. Fricker 2010; Mason 2015.

<sup>14</sup> Even if we are taken by Rosen’s argument, we might insist that certain kinds of morally grounded blame are appropriate, as long as they do not involve indignation (see e.g. Angela Smith’s (2013) understanding blame as a way of protesting—registering and challenging—a moral claim implicit in the blamee’s conduct), or think that angry reproach might be called for even when the targeted agents are not blameworthy (Calhoun 1989).

things as the slave's humanity or potential for happiness, or for the 1950s sexist to see that he is not taking his daughters' opportunities for flourishing seriously. If so, the shortcoming in quality of will required to explain their objectionable behavior might again be less pronounced than for their contemporary counterparts.

Another way in which EQW can accommodate intuitions about Rosen's cases relies on the fact that, unlike some ways of grounding blameworthiness in the agent's quality of will, EQW requires not only that the quality of will be bad. It also requires that the quality fall below what can be properly demanded, which in turn depends on the difficulty of caring, and so on the commitment needed to achieve and uphold various levels of caring. What matters here is that, typically, the commitment needed for achievement generally depends on external circumstances: the difficulty of achieving a certain language proficiency, for example, depends on how much of the language one can access and what type of feedback is available. Similarly, how easily I can sustain a concern for various morally relevant aspects of my actions might depend on the extent to which such a concern is supported by circumstances: it matters whether others share and voice their support for such concerns, and it might matter what sorts of opportunities I have to act on such concerns and to make them part of my identity (cf. Vargas 2013, Ch. 7–8). The circumstances of the Hittite lord seem to fall short in just this regard. Exactly because there is no societal support for the abolition of slavery, caring enough about these things to give up one's position as a slave owner would take extraordinary commitment or, in Rosen's words, a moral genius. Perhaps something similar can be said about the 1950s sexist. Given this, and given the possibility that less bad a quality of will is required to explain bad action supported by wrongheaded social mores, EQW at least opens up the possibility of withholding full or perhaps all moral blame from these agents.<sup>15</sup>

(It is worth noting that even if the ability of the *individual* 1950s sexist to care about gender equality is diminished because of his social circumstances, it might not have been difficult for the *group* of 1950s sexists to care a lot more. After all, if *they* had been committed to care a lot more about this, they would have found social support for their caring. This might ground shared responsibility, and shared obligations to change the circumstances (see Björnsson 2011; 2014a; Forthcoming-b).)

#### 4. Skeptical worries from explanatory innocence

Together with existing independent support for the Explanatory Quality of Will account, the considerations adduced in the previous section begin to undermine the requirement that blameworthiness be grounded in clear-eyed *akrasia*. Assuming that traditional sources of skepticism can be rejected, EQW seems to give the right verdict in cases meant to motivate the *akrasia* requirement, without overgeneralizing from these cases and implying that most of our positive attributions of blame are mistaken.<sup>16</sup> But I also believe that this and other quality

<sup>15</sup> Something similar can also be said about Kleinbart the Oblivious in Rosen 2008: 605–09.

<sup>16</sup> Even on these assumptions, it is quite possible that we frequently exaggerate blameworthiness by overestimating the amount of ill will required for a certain action. But the mistakes are likely less radical and pervasive than they would be given an absolute *akrasia* requirement.

of will-based responses to cases like Rosen's fail to address a deeper worry that has led a number of different people to independently make very similar arguments for epistemically based skepticism. More needs to be said to answer this worry.

Revisit the case of the 1950s sexist father. As Rosen tells his story, the father finds it obvious that men are entitled to better educational opportunities than women "because he was raised to find it obvious and because the people he takes seriously find it obvious". Perhaps, *given this background and this environment*, it is difficult for him to care as much about opportunities for his daughters as for his sons, and difficult to see that he gives his children unequal regard. We can add further details to the case, however, that make this seem less plausible. Suppose that there were people in his town, just down the street, who were known to support higher education for women. But suppose further that the father failed to take their perspective seriously because their overall political views were starkly at odds with views he was raised to find obvious and that people he took seriously found obvious: they were taken to be "commies", let us say.<sup>17</sup> Finally, suppose that if he had been somewhat committed to caring about giving his daughters the same opportunities as his sons, he would have begun talking to these people exactly because of their views about this matter. Given the support he would have received from this group, caring about equal opportunities would not have been very hard.

If this is how we understand the case, it might have been only marginally harder for the father to care equally about his daughters' and sons' opportunities than it is for someone today. Moreover, even if the father's judgment was in fact clouded by his social context, it would not have been clouded if he had cared enough about his daughters' opportunities to seek out a more beneficial social context. Given this, it is unclear how EQW can avoid the conclusion that the sexist father is quite blameworthy. But focusing on how he was in fact raised, and on the people he was raised to take seriously, moral indignation can still seem quite misplaced. Sure, it would not have been difficult for him to achieve or uphold an appropriate level of caring about educational opportunities for his daughters given a commitment to do so. But to form such a commitment in the first place under those circumstances would still have taken extraordinary moral insight and independence, and it seems wrong to blame someone for something that only a moral genius would have avoided.

I take cases like these to pose a real challenge to EQW. Or, rather, they pose a real challenge given DEMAND FROM CAPACITY, i.e. given that the degree of caring that can be properly demanded depends only on the difficulty of caring given circumstances and relevant cognitive and motivational capacities. One way of defending EQW would thus be to simply give up DEMAND FROM CAPACITY and say that nothing like the degree of caring that can be demanded of a father today could be demanded of Rosen's 1950s father, *even if caring to that degree back then would have been quite feasible*. On this interpretation, EQW would accommodate the sense that cultural context excuses. It would also likely have quite revisionary implications,

---

<sup>17</sup> In Rosen's own version of the story, the father was aware that there were people taking a different view of the matter, but because of his background, even reflecting "honestly and with something like an open mind (2003: 67)" on these views, they would strike him as wrong-headed.

as many a mistaken moral view and lack of moral caring will have a history similar to sexist outlook of the 1950s father.

There is reason to reject this response to the challenge, however. The kind of epistemically based skepticism that it would accommodate, I will suggest, is ultimately grounded in the same sort of intuition as traditional forms of skepticism. It thus poses no independent skeptical challenge to the combination of EQW and DEMAND FROM CAPACITY.

What seems to undermine blameworthiness in Rosen's cases is that the agent's behavior, although reflecting a poor quality of will, is *ultimately not his fault*, or due to any moral shortcoming of his. The worry is not exactly the traditional skeptical one that his actions might ultimately be explained by factors lying outside of his control: prior causes, external conditions, or chance. For all we have said, his lack of equal concern was not determined by factors outside of his control, and his control might not have been undermined by any indeterminacy involved. But on reflection, the epistemic and traditional skeptical worries alike seem to be grounded in the following thought:

EXPLANATORY INNOCENCE: The agent's bad behavior is ultimately explained by factors that do not ground blame.

Where traditional responsibility skeptics have stressed factors outside of the agent's control, epistemically based skepticism also focuses on agential aspects over which the agent might perhaps have had control but that do not constitute shortcomings: strategies of acting on one's internal reasons (Levy), or of assuming that one's moral outlook and those of people one trusts are basically sound (Rosen).

If this diagnosis is correct, I will argue, we can account for intuitions supporting both traditional and epistemically based responsibility skepticism while denying that they are veridical. This way of explaining away skeptical intuitions, defended for traditional forms of skepticism elsewhere (Björnsson and Persson 2012; 2013; Björnsson 2014b; 2016), builds on the following:

EXPLANATION EXPLANATION: Responsibility judgments involve explanatory judgments. Because of this, they vary with explanatory perspectives in ways characteristic for such judgments. Skeptical arguments induce non-standard explanatory perspectives that give rise to skeptical intuitions.

In what follows, I will first indicate how the EXPLANATION EXPLANATION can account for the appeal of traditional responsibility skepticism and incompatibilism, before returning to the epistemically based variety.

The general idea that judgments of responsibility, blame, and credit involve a species of explanatory judgment is not very controversial. It is often acknowledged or stressed that if an agent is morally blame for some event, it must have happened *because of her*, and we have seen how this can be seen as a general requirement on blame, whether moral or non-moral. More specifically, the assumption that responsibility judgments involve explanatory judgments can help explain the variety of excuses ordinarily acknowledged, and the role of partial justifications in mitigating blame (as sketched in section 2), along with otherwise puzzling


attributions of shared responsibility (See e.g. Björnsson 2011; Björnsson and Persson 2012; Björnsson Forthcoming-a). The key idea here, though, is that in virtue of being explanatory judgments, responsibility judgments are affected by the same kind of contextual features as everyday explanatory judgments.

The kind of context-dependence that concerns us is well known, and can be said to have two interrelated features.<sup>18</sup> First, ordinary explanatory judgments are strongly *selective*. If we are thinking about why some event E happened, we will focus only on one (or perhaps a few) events or conditions that were part of the causal prehistory of E, at the exclusion of innumerable others. If we are thinking about why a house burned down, for example, we might focus on the fact that the house was hit by lightning, rather than on the fact that the air contained oxygen, the fact that the house was built of combustible matter, or the fact that it lacked a first-class sprinkler system. Though we understand that these other factors were necessary conditions for E, we see them as part of the *background* against which something explains E rather than as *what explains E*, or its *significant explanation*. Our focus is typically on factors that deviate from what is generally expected, and so more informative than the factors relegated to the explanatory background. Moreover, we naturally focus on factors earlier in an explanatory chain, as long as they have a comparatively straightforward or familiar explanatory connection to E. Though we think that there was a series of intermediate events between the strike of lightning and the burning down of the house, such as the spread of fire in the attic just after the house was hit, we think of these exactly as *intermediary* events between what explanans and explanandum. Similarly, we acknowledge that the lightning that hit the house had a causal prehistory—the lightning was caused by a separation of charges in the neighboring atmosphere, say—but we focus on the lightning, as it has a more straightforward and familiar causal connection to the explanandum than have the events in its causal prehistory.

Schematically, our understanding of how multiple aspects of reality influence other aspects builds on a representation of how values of multiple variables are functionally related to values of other variables. In thinking that something explains something else, we (a) take explanandum and explanans as values of variables, (b) take the values of certain other variables as explanatory background, in effect treating them as constants, and (c) take the value of the explanandum variable to be a function of the value of the explanans variable given the values of the background variables, perhaps via some intermediary dependent variables. In doing so, we are ignoring various factors influencing the background constants as well as factors influencing the value of the explanans variable, treating it as an independent variable in our selective model of how the explanandum came to be. (See **Figure 1**.)

---

<sup>18</sup> My description of these in the next couple of paragraphs draws from Björnsson 2014b.

**Figure 1.** Explanatory model (grey area) with background conditions and independent, intermediary, and dependent variables

The second feature of the context-dependence of explanatory judgments is that the selection of something as the explanans variable depends on our explanatory interests and salient representations of the influence between variables. Though you and I might ordinarily focus on the lightning when thinking about why the house burned down, a fire engineer might instead focus on the lack of a lightning rod, thus relegating the fact that the house was hit by lightning to the explanatory background. Similarly, a politician thinking about the same event might focus on inadequate funding for the fire department, and a physicist on specific properties of the building materials. Because of different explanatory interests, they might relegate different factors to the explanatory background and focus on different relations of influence involving different variables. And because of this, they will then naturally think of different things as *the significant explanation of the outcome*.<sup>19</sup>

The assumption that responsibility judgments are selective and context-dependent along the lines of other kinds of explanatory judgments can help us explain a variety of otherwise disparate-looking facts about such judgments, or so I have argued (see e.g. Björnsson and Persson 2012; 2013; Björnsson 2014b). What matters for our purposes is that it provides a unified explanation of skeptical intuitions about responsibility, intuitions elicited by deterministic scenarios, regress arguments, and arguments from luck. Put very generally, the sort of account that I have defended starts with the idea that since we attribute moral responsibility based on seeing the agent's quality of will as the value of an independent variable, such attributions require the employment of certain kinds of folk-psychological explanatory models: call these "standard" models. What happens when we are confronted with descriptions of deterministic worlds or any of a wide family of skeptical arguments is that these trigger explanatory perspectives in which the agent's quality of will figures as a mere intermediary dependent variable (like the fire spreading in the attic). The standard model is then no longer front and center, and the agent no longer seems responsible.<sup>20</sup>

<sup>19</sup> The aspects of selectivity and context-dependence mentioned here do not exhaustively account for our causal-explanatory judgments. For my own take on the basic structure of our causal-explanatory judgments, see Björnsson 2007.

<sup>20</sup> Based on intriguing empirical evidence, Eddy Nahmias and Dylan Murray (2014) have proposed an alternative explanation of skeptical intuitions, suggesting that have as their source a confused understanding



For example, regress arguments such as Galen Strawson’s “basic argument” or typical descriptions of determinism trigger a move from the standard model to an abstract model where the agent’s quality of will fails to stand out as particularly significant among the events causally preceding the putative object of responsibility, instead making some first cause such as the Big Bang or the abstract category *prior causes* the independent variable (Björnsson and Persson 2012: 336–41).<sup>21</sup> Manipulation arguments work differently, instead triggering a move to a concrete model where one salient factor—the manipulation—straightforwardly explains the action by explaining the agent’s quality of will, thus rendering the latter a mere intermediate variable, undermining the sense that the agent is responsible.<sup>22</sup> In either case, once one has been taken by these new explanatory perspectives, they have their own *moral pull*: exactly because an agent seen from this perspective seems non-culpable, the perspective will seem to reveal the inappropriateness of retributive emotions and actions. Consequently, changing focus back to the standard perspective and ignoring prior causes would seemingly ignore something of great moral importance.

This EXPLANATION EXPLANATION of forms of skepticism makes sense of variations in intuitions supporting traditional skeptical or incompatibilist worries, among philosophers and laymen alike (for data on the latter, see, among many others, Nahmias et al. 2007; Nichols and Knobe 2007). Unlike explanations of traditional skeptical worries in terms of incompatibilist conceptual requirements on responsibility, the EXPLANATION EXPLANATION can make straightforward sense of compatibilist intuitions too, in terms of judgments made from standard explanatory models. Indeed, given that standard models are central to social cognition and given that attributions of moral blame and credit play central roles in moral thinking and discourse, explanatory perspectives supporting such attributions will have their own attraction, counteracting that of non-standard skeptical perspectives. Given the EXPLANATION EXPLANATION, the attraction of both skeptical and non-skeptical perspectives is to be expected.

The suggestion now is that cases adduced in support of epistemically based skepticism work similarly to manipulation cases. The cases motivating either of these forms of skepticism focus our attention on factors that explain the action in question without invoking any

---

of determinism and its implications for agency. This suggestion, however, is undermined by newer studies (Rose and Nichols 2013; Björnsson 2014b).

<sup>21</sup> This, I have argued, is also why deterministic worlds do much less to undermine ascriptions of responsibility when people consider concrete acts of wrongdoing than when considering the abstract question of whether agents can be fully responsible in that world (Björnsson and Persson 2013; Björnsson 2014b).

<sup>22</sup> Many arguments from luck or disappearing agency work differently still. By focusing on some relevant (random) element not part of the agent’s quality of will, they in effect relegate the latter to the explanatory background, much as the fire engineer’s explanatory interests did the lightning (see Björnsson and Persson 2012: 341–44).

Chandra Sripada (2012) takes studies of his to suggest that people take manipulation to undermine responsibility because they take it to undermine certain compatibilist conditions. However, further studies undermine Sripada’s interpretation of the data and provide contrary evidence in line with the current proposal (Björnsson 2016).

substandard quality of will of the agent. In doing so, they prompt us to see his quality of will as, at best, a mere dependent variable. For example, in considering the behavior of Rosen's sexist father in light of his upbringing and social surroundings, it is natural to see one or more non-culpable factors (his normal capacities for moral reasoning, his upbringing, his social surroundings) as independent variables, while seeing his lack of concern for the opportunities of his daughters as a mere intermediary dependent variable. Insofar as we keep seeing the situation from this perspective rather than from a standard perspective, the father will not strike us as morally to blame.

The non-standard perspective introduced also has its own attraction. In manipulation cases, the manipulator that controls the agent's action by controlling her quality of will stands out from what is ordinarily expected and makes salient a non-standard explanatory model in which the agent's quality of will is seen as a mere dependent variable. In the case of the sexist father, the fact that his bad quality of will is the result of seemingly non-culpable epistemic strategies is itself remarkable and provides salience to the corresponding non-standard explanatory model. In either case, however, the sense that the morally ignorant agent's responsibility is undermined can be resisted by resisting the accompanying explanatory perspective and stressing the involvement of the agent's substandard quality of will (see e.g. Harman 2011; Wieland 2015). For example, if we focus on the fact that his daughters missed out on an education because their father cared more about the opportunities of his sons, without attending to the causal prehistory of his unequal caring, some degree of indignation on part of the daughters seems appropriate (cf. Rosen 2003: 73).

The EXPLANATION EXPLANATION of traditional and epistemically based skepticism does not in itself show that these forms of skepticism are mistaken. In fact, if the explanation is correct, it should ring true to those having skeptical intuitions. But if it is correct, this suggests that intuitions about cases cannot tell us whether to accept the skeptical conclusions until we know whether responsibility judgments are more reliable when made from standard or from non-standard perspectives (cf. Björnsson and Persson 2012: 345–48). In themselves, skeptical intuitions thus constitute no direct threat to the combination of EQW and DEMAND FROM CAPACITY. Moreover, assuming that we have reason to reject intuitions about responsibility stemming from non-standard explanatory perspectives, epistemically based skepticism motivated by cases like Rosen's can be rejected along with more traditional, metaphysically based forms. And if we have independent reason to accept EQW and DEMAND FROM CAPACITY, the EXPLANATION EXPLANATION lets us explain away the intuitions seemingly supporting all these forms of skepticism.

### **Acknowledgments**

This paper has benefited from comments from audiences at Umeå University, the University of Gothenburg, and the 2015 New Orleans Workshop on Agency and Responsibility, from two anonymous referees for Oxford Studies in Agency and Responsibility, and from Zoë Johnson King and Andreas Brekke Carlsson, who each provided extensive written and very

thoughtful feedback. Work on this paper was funded by Riksbankens Jubileumsfond and the Swedish Research Council.

### Bibliography

- Arpaly, Nomy 2003: *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford University Press.
- Arpaly, Nomy 2006: *Merit, Meaning, and Human Bondage: An Essay on Free Will*. Princeton, NJ: Princeton University Press.
- Arpaly, Nomy and Schroeder, Timothy 2014: *In Praise of Desire*. Oxford University Press.
- Björnsson, Gunnar 2007: 'How Effects Depend on Their Causes, Why Causal Transitivity Fails, and Why We Care About Causation'. *Philosophical Studies*, 133, pp. 349-90.
- Björnsson, Gunnar 2011: 'Joint Responsibility without Individual Control: Applying the Explanation Hypothesis'. In *Moral Responsibility: Beyond Free Will and Determinism*. Vincent, Nicole, van de Poel, Ibo and van den Hoven, Jeroen (eds) Dordrecht: Springer pp. 181–99.
- Björnsson, Gunnar 2014a: 'Essentially Shared Obligations'. *Midwest Studies in Philosophy*, 38, pp. 103–20.
- Björnsson, Gunnar 2014b: 'Incompatibilism and 'Bypassed' Agency'. In *Surrounding Free Will*. Mele, Alfred (ed) New York: Oxford University Press pp. 95–122.
- Björnsson, Gunnar 2016: 'Outsourcing the Deep Self: Deep Self Discordance Does Not Explain Away Intuitions in Manipulation Arguments'. *Philosophical Psychology*.
- Björnsson, Gunnar Forthcoming-a: 'Explaining (Away) the Epistemic Condition on Moral Responsibility'. In *Responsibility—the Epistemic Condition*. Robichaud, Philip and Wieland, Jan Willem (eds) New York: Oxford University Press.
- Björnsson, Gunnar Forthcoming-b: 'Individual and Shared Obligations: In Defense of the Activist's Perspective'. In *Philosophy and Climate Change*. Budolfson, Mark, McPherson, Tristram and Plunkett, David (eds) Oxford University Press.
- Björnsson, Gunnar and Pereboom, Derk 2016: 'Traditional and Experimental Approaches to Free Will and Moral Responsibility'. In *Companion to Experimental Philosophy*. Sytsma, Justin and Buckwalter, Wesley (eds) Wiley pp. 142–57.
- Björnsson, Gunnar and Persson, Karl 2012: 'The Explanatory Component of Moral Responsibility'. *Noûs*, 46, pp. 326–54.
- Björnsson, Gunnar and Persson, Karl 2013: 'A Unified Empirical Account of Responsibility Judgments'. *Philosophy and Phenomenological Research*, 87, pp. 611–39.
- Calhoun, Cheshire 1989: 'Responsibility and Reproach'. *Ethics*, 99, pp. 389-406.
- Clarke, Randolph 2014: 'Negligent Action and Unwitting Omission'. In *Surrounding Free Will*. New York: Oxford University Press pp. 298–317.
- Fischer, John Martin and Ravizza, Mark 1998: *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge U. P.
- FitzPatrick, William J. 2008: 'Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge'. *Ethics*, 118, pp. 589-613.
- Frankfurt, Harry G. 1969: 'Alternate Possibilities and Moral Responsibility'. *Journal of Philosophy*, 66, pp. 829-39.
- Fricker, Miranda 2010: 'I—the Relativism of Blame and Williams's Relativism of Distance'. *Aristotelian Society Supplementary Volume*, 84, pp. 151-77.
- Harman, Elizabeth 2011: 'Does Moral Ignorance Exculpate?'. *Ratio*, 24, pp. 443-68.
- Levy, Neil 2009: 'Culpable Ignorance and Moral Responsibility: A Reply to Fitzpatrick'. *Ethics*, 119, pp. 729-41.
- Levy, Neil 2011: *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. Oxford University Press.
- Mason, Elinor 2015: 'Moral Ignorance and Blameworthiness'. *Philosophical Studies*, pp. 1-21.
- Massoud, Amy 2016: 'Moral Worth and Supererogation'. *Ethics*, 126, pp. 690-710.
- McKenna, Michael 2012: *Conversation and Responsibility*. New York: Oxford University Press.
- Murray, Dylan and Nahmias, Eddy 2014: 'Explaining Away Incompatibilist Intuitions'. *Philosophy and Phenomenological Research*, 88, pp. 434-67.

- Nahmias, Eddy, Coates, D. Justin and Kvaran, Trevor 2007: 'Free Will, Moral Responsibility, and Mechanism: Experiments on Folk Intuitions'. *Midwest Studies in Philosophy*, 31, pp. 214-42.
- Nichols, Shaun and Knobe, Joshua 2007: 'Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions'. *Noûs*, 41, pp. 663-85.
- Robichaud, Philip 2014: 'On Culpable Ignorance and Akrasia'. *Ethics*, 125, pp. 137-51.
- Rose, David and Nichols, Shaun 2013: 'The Lesson of Bypassing'. *Review of Philosophy and Psychology*, 4, pp. 599-619.
- Rosen, Gideon 2003: 'Culpability and Ignorance'. *Proceedings of the Aristotelian Society*, 103, pp. 61-84.
- Rosen, Gideon 2004: 'Skepticism About Moral Responsibility'. *Philosophical Perspectives*, 18, pp. 295-313.
- Rosen, Gideon 2008: 'Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility'. *The Journal of Philosophy*, 105, pp. 591-610.
- Sher, George 2009: *Who Knew?* Oxford U. P.
- Shoemaker, David 2013: 'Qualities of Will'. *Social Philosophy and Policy*, 30, pp. 95-120.
- Smith, Angela 2013: 'Moral Blame and Moral Protest'. In *Blame: Its Nature and Norms*. Coates, D. Justin and Tognazzini, Neal A. (eds) Oxford University Press pp. 27-48.
- Smith, Angela M. 2005: 'Responsibility for Attitudes: Activity and Passivity in Mental Life'. *Ethics*, 115, pp. 236-71.
- Smith, Holly M. 2011: 'Non-Tracing Cases of Culpable Ignorance'. *Criminal Law and Philosophy*, 5, pp. 115-46.
- Sripada, Chandra Sekhar 2012: 'What Makes a Manipulated Agent Unfree?'. *Philosophy and Phenomenological Research*, 85, pp. 563-93.
- Strawson, Peter F. 1962: 'Freedom and Resentment'. *Proceedings of the British Academy*, 48, pp. 187-211.
- Vargas, Manuel 2013: *Building Better Beings: A Theory of Moral Responsibility*. Oxford: Oxford University Press.
- Wieland, Jan Willem 2015: 'What's Special About Moral Ignorance?'. *Ratio*, pp. n/a-n/a.
- Williams, B. A. O. 1976: 'Moral Luck'. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 50, pp. 115-35.
- Zimmerman, Michael J. 1997: 'Moral Responsibility and Ignorance'. *Ethics*, 107, pp. 410-26.
- Zimmerman, Michael J. 2008: *Living with Uncertainty*. Cambridge U. P.