

Derivative culpability

Martin Montminy

Department of Philosophy, University of Oklahoma, Norman, OK, USA

ABSTRACT

I explore the question of when an agent is *derivatively*, rather than *directly*, culpable for an undesirable outcome. The undesirable outcome might be a harmful incompetent or unwitting act, or it might be a harmful event. By examining various cases, I develop a sophisticated account of indirect culpability that is neutral about controversies regarding normative ethical issues and the condition on direct culpability.


ARTICLE HISTORY Received 22 May 2017; Accepted 12 February 2018

KEYWORDS Moral responsibility; indirect blameworthiness; culpable ignorance; causation

1. Introduction

Under what conditions is an agent *derivatively*, or *indirectly*, morally culpable for an undesirable outcome? Intuitively, three conditions must hold: the undesirable outcome must causally result from an earlier act performed by the agent; that earlier act must be morally wrong; and the agent must be directly culpable for that earlier act. As I will show, these three conditions are much too rough. I will examine a slightly more sophisticated account proposed by Holly Smith, and show that it faces several counterexamples. The three conditions on derivative culpability, I will argue, should be more specific and exhibit a more intimate relationship with each other than is usually thought. My account of derivative culpability will remain neutral about two separate, but related questions, namely ‘What, exactly, makes an action morally wrong?’ and ‘Under what conditions is an agent *directly* culpable for a wrongdoing?’

Before I proceed, I should note the connection between the topic of this essay and the notion of *tracing*: as many have pointed out, an agent is derivatively culpable for an outcome only if this outcome can be traced to an earlier (directly) culpable act. Hence, an account of indirect culpability is automatically

CONTACT Martin Montminy  montminy@ou.edu

© 2018 Canadian Journal of Philosophy

an account of tracing, when it concerns culpability. In Section 6, I will examine a recent debate about tracing that bears some relevance to the topic of this essay.¹

2. Three paradigmatic cases

To help fix ideas, let us consider three paradigmatic instances of derivative culpability. The first one involves an agent who culpably does something that reduces his abilities and later fails to do the right thing because of his incompetence.

Drink and Drive. Drunk Driver drinks more than he should, fully aware that he will later have to drive home while drunk. On his way home, he runs over a pedestrian. Drunk Driver would have had no trouble avoiding the accident had his driving abilities not been seriously impaired by alcohol.

Drunk Driver's culpability for harming the pedestrian is entirely traceable to his earlier act of drinking.² Once he is behind the wheel, it is beyond Drunk Driver's reduced abilities to do the right thing. Hence, because of his incompetence, he is derivatively rather than directly blameworthy for the accident. (I use 'blameworthy' and 'culpable' interchangeably throughout this paper.) Here, the agent's impaired abilities are, mostly, his motor skills. But similar cases of derivative culpability could involve reduced cognitive, motivational or perceptual abilities.

In the introductory section, I wrote that my account will be neutral about the conditions on direct culpability. However, I should specify that I treat the question whether an agent is blameworthy for a wrongdoing as distinct from the question whether a sanction or punishment would be appropriate for that wrongdoing. This position is relatively uncontroversial. Even authors who hold that blameworthiness is ultimately analyzable in terms of the attitude of holding responsible or blaming are careful to distinguish that attitude from moral sanction.³ The current discussion thus merely concerns the question whether Drunk Driver is blameworthy for harming the pedestrian. It is silent on what punishment he should receive, and more generally on how we should *overtly* respond to his wrongdoing.

Next, consider a case of *culpable ignorance*. In the following case described by Philip Robichaud and Jan Willem Wieland ([forthcoming](#)), the agent's original culpable act compromises not her abilities, but her knowledge of relevant facts.

Wrong Prescription. Dr. Lazy is aware that in her specific area she is morally required to spend approximately 10 h per week keeping up on practice-relevant research, and she heard that a colleague's recent article reports new, important findings about the drug Inscientium. Dr. Lazy has the time and energy to read it during work hours, but instead chooses to have coffee with a colleague. She knows that she should read the study, but opts for the coffee break out of frustration over how far behind she is on her research. In fact, in the past week she did not spend *any* time reading practice-relevant journals. In the study that Dr. Lazy fails to read, Inscientium – the most effective drug for treating hay fever – is decisively shown to cause fatal heart attacks in people with rare kidney conditions. The next day, Dr. Lazy prescribes Inscientium to a patient with one of the rare kidney conditions

who suffers a fatal heart attack as a result. (Robichaud and Wieland forthcoming, slightly edited.)

As Robichaud and Wieland acknowledge, their stipulation that a doctor ought to spend 10 h per week reading relevant journals is somewhat arbitrary. However, this level of precision is unimportant in this case. What matters is that Dr. Lazy has some obligation to inform herself about the relevant literature in her field. This obligation, Robichaud and Wieland remind us, stems from the more general obligation we all have to inform ourselves about the morality of our future conduct. Gideon Rosen coins the useful phrase ‘procedural epistemic obligation’ to designate this phenomenon: ‘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*’ (Rosen 2004, 301). Now, given that she is aware of the existence of a recent article that contains information about a drug she frequently prescribes, Dr. Lazy is very plausibly morally obligated to read the article in question, and is blameworthy for failing to meet this procedural obligation. Robichaud and Wieland’s intuition, which I share, is that Doctor Lazy is indirectly blameworthy for prescribing Inscientium, because she is directly blameworthy for not reading the study about that drug.

Wrong Prescription raises another issue, though. As Holly Smith remarks, ‘There are many cases in which enquiry should be made earlier, but it is better to enquire now rather than act without its benefit. In such a case the agent is indeed culpable for his act, but the culpability might be wholly traceable to the fact that he knows himself to be performing an act less good than its alternative, namely conducting further enquiry. Hence we cannot unhesitatingly assign all the culpability – or even any of it – to the fact that he is culpably ignorant when he acts’ (Smith 1983, 546). One could thus argue that Dr. Lazy is *directly* blameworthy for prescribing Inscientium: given its rumored important findings concerning the drug she is about to prescribe, she should read the recent journal article before she prescribes the drug. To avoid this problem, I will follow Smith and stipulate that ‘it is false that the agent ought now to enquire further before doing the act in question – either because he is unable to make further enquiry, or because he ought not to do so (perhaps the costs of enquiry, in terms of opportunities lost, are too great at this point)’ (ibid., 546–547). We should thus assume that is too late for Dr. Lazy to read the journal article: she must decide *now* how to treat her patient. Hence, her culpability for her unwitting act is entirely attributable to her earlier omission.

Our third case involves only one act, a directly culpable act that has harmful consequences.

Forest Fire. Breaking camp on a windy day, Hiker leaves a campfire burning. Sparks spread, and a forest fire results.⁴

In the first two cases, agents are derivatively to blame for actions they perform. These actions are bad consequences of earlier actions (or omissions). The agents are also indirectly blameworthy for consequences that are not themselves actions. In *Wrong Prescription*, Doctor Lazy is indirectly blameworthy not only for prescribing Inscientium to her patient, but also for the patient's heart attack. Similarly, in *Forest Fire*, Hiker is indirectly blameworthy for the forest fire, a bad consequence of his omission. It is generally admitted that one may be directly blameworthy only for actions or omissions.⁵ I thus endorse Smith's (1983) proposal to treat culpability for the harmful events caused by actions in the same way we do actions for which agents are indirectly blameworthy.⁶ Intuitively, this makes sense, since Hiker's culpability for the forest fire seems entirely due to his culpability for failing to douse the campfire. Moreover, as I will try to show in this paper, it is possible to offer a uniform treatment of all of these forms of derivative culpability.

3. Smith's account

Derivative culpability must be traceable to an earlier act. Following Holly Smith (1983), let us call that earlier act a *benighting act*, with the understanding that omissions count as acts. A quick glance at the three cases described in the previous section suggests that three conditions must hold for an agent to be derivatively culpable. First, the harm for which the agent is derivatively culpable must be caused by her benighting act. Plausibly, an agent is not to blame for an outcome if none of what she did (or failed to do) was a cause of that outcome. In the three cases, the undesirable outcome (i.e. incompetent act, unwitting act or a harmful event) is assumed to be a causal consequence of the benighting act. In *Wrong Prescription*, for example, it is assumed that Dr. Lazy would *not* prescribe Inscientium to her patient if she had read the study and learned about its potential lethal effect. Hence, her omission (and ignorance) causally contributes to her act. Let us call this the *causal condition*. Second, the benighting act must be morally impermissible. A morally permissible action may have some bad consequences. A beneficial medical treatment may cause some pain or discomfort, for example. But surely, a doctor is not to blame for the pain resulting from that treatment. Hence, very plausibly, the agent's benighting act should be morally wrong.⁷ Let us call this the *wrongness condition*. Third, the agent should be directly culpable for the benighting act. If Hiker were not culpable for not putting out the campfire, because, say, he was incapable of doing so (through no fault of his), then he would not be to blame for the forest fire. Let us call this the *culpability condition*. This condition echoes the point made by many philosophers, according to which derivative blameworthiness must entirely trace back to direct blameworthiness.⁸

The following *simple account* combines these three conditions: an agent is derivatively culpable for an undesirable outcome (i.e. an incompetent act, an unwitting act or a harmful event) just in case (1) the agent performed a morally wrong benighting act, (2) the benighting act gave rise to the undesirable outcome, and (3) the agent is directly culpable for the benighting act.

Unfortunately, the simple account is inadequate. Consider the following counterexample, adapted from (Smith 1983, 550):

Blood Type. To assist his son in his school project, Prof. Gamer promises to survey the members of his department about their blood types, and send him the results by email at noon. Instead of fulfilling his promise, Prof. Gamer plays videogames on his computer all morning. Early in the afternoon, a colleague suffers a serious accident causing major bleeding. Paramedics arrive and could start a transfusion immediately if the colleague's blood type were known. Unfortunately, because she is unconscious, she cannot provide this information, and dies before the transfusion can be performed. Had Prof. Gamer fulfilled his promise, he would have known his colleague's blood type, and her life would have been saved.⁹

Prof. Gamer's failure to obtain the information is both wrong and culpable, and it led to his inability to help save his colleague's life. However, Prof. Gamer is not (derivatively) to blame for the colleague's death.

Smith's own account avoids this problem.¹⁰ She is primarily concerned with cases of culpable ignorance, where an agent's benighting act gives rise to an unwitting 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' (ibid., 566). Smith discusses other forms of derivative culpability, and explains Hiker's derivative culpability in *Forest Fire* very similarly: 'We say the hiker is morally to blame for the forest fire. We are willing to make such a judgment when three conditions are met: (1) the person is to blame for performing some action (here, failing to douse the campfire), (2) the action gives rise to an undesirable event (here, the forest fire), and (3) the person knows the action risks giving rise to that event' (ibid., 564).

Smith's account differs from the simple account in two important respects. First, there is no wrongness condition in Smith's account. However, she clearly accepts this condition (ibid., 547–548). Second, Smith further specifies the culpability condition by including an *epistemic requirement*: the agent knew that his benighting act risked giving rise to a negative outcome. This epistemic requirement is not satisfied in *Blood Type*: Prof. Gamer did not know that his failure to survey his colleagues risked leading to his inability to help save one of them. As Smith puts it, the undesirable outcome did not 'fall within the known risk of the benighting act' (ibid., 551).¹¹

For future reference, let us restate Smith's account, adding a wrongness condition, replacing 'unwitting act' by 'undesirable outcome' for the sake of generality, and incorporating the epistemic requirement into the culpability condition. Smith holds that an agent is derivatively culpable for an undesirable outcome

just in case (1) the agent performed a morally wrong benighting act, (2) the benighting act gave rise to the undesirable outcome, and (3) the agent is directly culpable for performing the benighting act, and while performing that act the agent knew that it risked giving rise to the undesirable outcome. Unfortunately, this account is still too rough. I will examine each condition in turn.

4. The wrongness condition

According to the wrongness condition, the benighting act giving rise to the undesirable outcome must be morally impermissible. But this condition is too weak.

Extended Stay. Slacker Dad is with his three-year-old daughter at the park, and has promised his wife that they would be back home at 4:30 pm, to help with dinner preparation and chores. But right before it is time to head back home, his daughter's good friend arrives at the park. Slacker Dad decides to extend their stay. He is aware that despite their occasional fights, his daughter and her friend get along well and much enjoy each other's company. Unfortunately, a few minutes later, the two girls get into an argument that ends in tears.

Smith's three conditions are satisfied here. Given the promise made to his wife, it was wrong for Slacker Dad to extend their stay at the park. The extended stay gave rise to the girls' fight: no fight would have occurred if they had left at the promised time. Moreover, Slacker Dad is to blame for extending their stay at the park, since he is aware that he is breaking his promise to his wife. Finally, Slacker Dad knew that staying at the park could lead to a fight. But, intuitively, Slacker Dad is not to blame for the fight.

Interestingly, some passages in Smith's essay can help address the problem raised by *Extended Stay*. In her discussion of *Blood Type*, Smith (ibid., 551) briefly mentions a *second* reason why Prof. Gamer is not derivatively culpable for his colleague's death: what makes his benighting act wrong is not the fact that it risked making him unable to help save his colleague, but the fact that he was breaking his promise to his son. The wrongness condition should thus be amended: the fact that it risked giving rise to the undesirable outcome must be what makes the benighting act wrong.¹² This condition is not satisfied in *Extended Stay*, for the fact that it risks giving rise to a fight does not make Slacker Dad's extended stay wrong. Even though they occasionally fight, it is still worth it for the two girls to spend time together. What makes Slacker Dad's act wrong is the fact that it breaks his promise to his wife.

Despite appearances, this new wrongness condition should not be interpreted as being committed to consequentialism. Consider how this condition would apply to *Wrong Prescription*. A Kantian, for example, could hold that what makes Dr. Lazy's omitting to read the journal article morally wrong is the fact that it risks failing to respect her patients. And a Ross-style theorist could contend that her omission is wrong because it risks violating a duty of non-maleficence,

that is, a duty not to harm others. On the assumption that either one of these assessments is correct, the wrongness condition should be understood as met.

The new wrongness condition is slightly too strong, though.

Multiple Side Effects. A certain medication M has three distinct potential side effects. If it had only one of the three, the treatment would be worth it, given its almost guaranteed benefit. However, when combined, the three potential side effects outweigh the potential benefit. Fortunately, there exists alternative medications that are just as effective as M and have only one of the three potential side effects M has. Despite being aware of that, Dr. Reckless prescribes medication M, and the following day, the patient suffers from one – but only one – of the side effects.

Dr. Reckless is directly to blame for his bad prescription, and derivatively to blame for the side effect. However, the fact that his action risked causing that side effect did not *by itself* make his action wrong; it merely *contributed* to its wrongness. The wrongness condition should thus be weakened: the fact that it risked giving rise to the undesirable outcome must contribute to the wrongness of the benighting act. I should note that the locution ‘contributes to the wrongness of the act’ is understood to entail that the act is morally wrong. A doctor is not to blame for a side effect resulting from a treatment whose benefits clearly outweigh its potential side effects. According to the proposed terminology, this side effect would *not* be wrong-contributing

5. The causal condition

According to the causal condition, the benighting act must give rise to the undesirable outcome. Some may reject the idea that culpability for an outcome requires the benighting act’s *causing* the outcome. In *Forest Fire*, Hiker’s culpability for the forest fire is due not to something he *did*, but to something he *failed* to do, namely put out the campfire. But it is controversial whether omissions such as Hiker’s can be causes. An omission might be thought of as an absence or non-occurrence, and some authors deny that absences or non-occurrences are causally efficacious.¹³ Fortunately, accounts of derivative culpability can steer clear of that controversy. Philosophers who hold that omissions, and more generally absences, cannot be causes usually grant that they may be *quasi-causes*.¹⁴ Quasi-causation can be understood counterfactually. Hiker’s omission is a quasi-cause of the forest fire, since, roughly, his putting out the campfire *would* have interfered with an actual causal process that led to the forest fire.¹⁵ For simplicity’s sake, I will now use ‘gives rise to,’ ‘causes’ and related terms to designate causes or, if there are such things, quasi-causes.

The causal condition appears to need strengthening, since the causal path connecting the benighting act and the undesirable outcome may be *deviant*.¹⁶

Questionable Lesson. Breaking camp on a windy day, Hiker leaves a campfire burning. Luckily, the campfire is soon extinguished by the wind. However, Hardnosed, who wants to teach Hiker a lesson, was watching the whole scene. Unbeknownst

to Hiker, Hardnosed reignites the campfire, which eventually spreads to its surrounding and produces a forest fire.

A second agent is of course not necessary to generate a deviant causal chain. We could instead imagine that after Hiker leaves a campfire burning, the heat emanating from it startles a grazing buffalo. The buffalo runs away and causes an enormous buffalo herd stampede. The stampede then triggers a debris avalanche that releases a lava flow from a nearby volcano, which eventually produces a forest fire. But from now on, I will limit the discussion to the simpler *Questionable Lesson*.

In *Questionable Lesson*, Hiker's failure to put out the campfire gave rise to the forest fire: Hardnosed would not have done what he did had he seen Hiker dousing the campfire. Furthermore, the wrongness and culpability conditions are also satisfied. But although he is blameworthy for not putting out the campfire, Hiker is clearly not to blame for the forest fire.¹⁷ The causal condition, it seems, needs to be revised: the benighting act must give rise to the undesirable outcome through a *non-deviant* causal path.

However, I will propose a different solution to the problem raised by *Questionable Lesson*. Consider the question of what makes a causal path *deviant*. Plausibly, a deviant causal path is one that does not affect the moral assessment of the action that initiated it. Intuitively, the fact that Hardnosed reignited the campfire, which eventually led to a forest fire, does *not* contribute to the wrongness of Hiker's omission. Most normative ethical theories would entail that. But not all. A brief discussion of a debate among consequentialists will help appreciate this point.¹⁸ According to actual consequentialism, the morality of an act is determined by its *actual* consequences. An action is morally impermissible, even if it is a matter freak accident, unpredictable in advance by anybody, that it has bad consequences. And the fact that, given available evidence, it was likely to have catastrophic consequences does not affect the morality of an act negatively if these consequences do not obtain. On this view, non-deviant causal paths are simply *actual* ones. Hence, on this view, the causal path leading to the forest fire in *Questionable Lesson* is non-deviant.

Expected consequentialism, on the other hand, holds that only *probable* (or *foreseeable*) consequences matter to morality. Expected consequentialism is really a family of views that differ in how *probable* consequences are characterized. Should we consider only the evidence actually available to the agent, or should we also include the evidence she would have access to in nearby worlds? Proponents of expected consequentialism may answer this question differently. But they agree on the generic condition that a causal path is non-deviant if it is probable, given a relevant body of evidence. According to expected consequentialism, the causal path leading to the forest fire in *Questionable Lesson* is deviant, since it is not expectable. Hence, what makes Hiker's failure to put out the campfire wrong is not the fact that it actually gave rise to the forest fire, but

the fact that there were probable (but unrealized) causal paths going from the unextinguished campfire to a forest fire.

These considerations show how to accommodate cases such as *Questionable Lesson*, without invoking the notion of non-deviant causal paths explicitly. Let us keep the causal condition as originally stated, but adjust the wrongness condition as follows: the fact that the benighting act risked giving rise to the undesirable outcome in the way it actually did must contribute to the wrongness of that act. Our two brands of consequentialism would both take the causal condition to be satisfied in *Questionable Lesson*, since Hiker's omission did give rise to the Forest Fire. What about the wrongness condition? The two views would differ here. Expected consequentialism would hold that the fact that Hiker's failure to put out the campfire risked giving rise to forest fire in the way it actually did does not contribute to the wrongness of his omission. This is because *given available evidence*, that particular risk was negligible. For this reason, the causal path leading to the forest fire is deviant, and hence does not matter to the moral assessment of Hiker's omission. The latter is morally wrong, not because it led to the forest fire in the way it actually did, but because the available evidence entailed that there were many ways his omission could have caused the forest fire. By contrast, actual consequentialism would hold that the fact that Hiker's failure to put out the campfire risked giving rise to forest fire in the way it actually did does contribute to the wrongness of his omission. This is because *given what actually happened*, the risk in question was realized.

Now, although they hold that both the causal and wrongness conditions are satisfied in *Questionable Lesson*, actual consequentialists need not deem Hiker blameworthy for the forest fire. They can insist that the culpability condition is not satisfied in this case, for Hiker did not know that a forest fire risked being caused in the way it actually did. This means that the culpability condition must be revised in a similar way the wrongness condition was: the agent is culpable for performing the benighting act, and while performing that act the agent knew that it risked giving rise to the undesirable outcome in the way it actually did.

Once again, I should note that this issue is not of concern only for consequentialists. The dispute between actual and expected consequentialists I just presented is part of a larger debate between objective and subjective moral theories. According to purely subjective moral theories, the moral rightness of an action is fixed by the agent's beliefs or evidence about the circumstances (and consequences) of the action.¹⁹ Such theories are, of course, controversial.²⁰ There is a wide range of alternatives. At the other end of the spectrum, we find radical objectivist theories, according to which the morality of an action depends solely on the objective features of the agent's circumstances. Actual consequentialism is a version of this view. To navigate between radical subjectivism and radical objectivism, we may allow epistemic constraints on what makes a descriptive fact relevant to the moral rightness of an action. A relevant descriptive fact need not be actually known by the agent; however, it should

be *in some sense* epistemically accessible to her. For example, consider a doctor who genuinely believes that the treatment she is prescribing for a patient will be beneficial. Proponents of a moderate objective theory may hold that it is morally wrong for the doctor to do so on the grounds that it is relatively well known within her specialization that the treatment is harmful. However, they may add, on the assumption that her ignorance is not itself culpable (because she fulfilled her procedural epistemic obligations), she would not be to blame for the harm caused by the treatment. This is another illustration of how a more objective moral view can deem an agent blameless for an act (or an undesirable outcome), despite the fact that the act (or the causing of that undesirable outcome) was wrong.

6. Taking stock

Here is our *revised account*: an agent is indirectly to blame for an undesirable outcome just in case (1) the agent performed a benighting act that is wrong at least in part because it risked giving rise to the undesirable outcome in the way it actually did, (2) the benighting act gave rise to the undesirable outcome, and (3) the agent is directly culpable for the benighting act, and while performing that act, she knew that it risked giving rise to the undesirable outcome in the way it actually did.

The revised account will require one more revision. Before I explain why, I would like to address two potential concerns. First, one might wonder whether *each* of the three conditions of the revised account is required. Couldn't we dispense with, say, the wrongness condition? No. To appreciate this point, let us consider Robichaud and Wieland's ([forthcoming](#)) discussion of indirect culpability. Robichaud and Wieland are primarily interested in instances of culpable ignorance. By examining a series of cases, they consider an account akin to the revised account, and raise an objection against it. More specifically, their objection concerns a wrongness condition similar to the one included in the revised account. According to what they call the *basing constraint*, the reason why the benighting act is wrong must match the reason why the subsequent unwitting act is wrong. Such a match clearly obtains in *Wrong Prescription*, for instance. It is wrong for Dr. Lazy to prescribe Inscientium because this drug causes avoidable harm, and it is wrong for Dr. Lazy not to read the study about the drug for precisely this reason. Setting aside the refinements prompted by considerations about *Multiple Side Effects* and *Questionable Lesson*, the basing constraint is very close to my wrongness condition. Robichaud and Wieland hold that the following case 'makes more decisive trouble for the basing constraint' (*ibid.*):

Drug-Induced Deafness. Dr. Negligent thinks that she should read a new article on the dangers of the drug Surditas because it is rumored to cause temporary and mild hearing loss. Although she has not read the required amount of practice-relevant research in the past week and has time and energy to read the article, she

decides to have coffee with a colleague. Because she fails to read the article she does not learn that the drug actually causes permanent deafness in patients with rare throat conditions rather than the milder symptoms that were rumored. The next day, Dr. Negligent prescribes the drug to a patient with this rare condition who subsequently goes deaf. If the rumor was that Surditas would cause total deafness, she definitely would have read it. We know this, because a week earlier, she read a different article about a different drug precisely because it was rumored to cause total deafness. (ibid., slightly edited)

Clearly, the causal condition is respected in this case, since Dr. Negligent's failure to read the article is a cause of her unwittingly harmful prescription. Is the basing constraint satisfied? Yes. As Robichaud and Wieland remark, 'not reading the article is wrong because Surditas causes permanent deafness and reading the study would have enabled her to see that she should not prescribe this drug' (ibid.). The reason Dr. Negligent's omission is wrong thus matches the reason her subsequent unwitting act is wrong. (The wrongness condition is also satisfied: Dr. Negligent's omission is wrong because it risked giving rise to the patient's deafness in the way it actually did).²¹ Finally, Robichaud and Wieland point out, her awareness that she should read the article makes it plausible that Dr. Negligent is blameworthy for not doing so. The problem, they write, is that intuitively, Dr. Negligent is not blameworthy for her patient's deafness. This is because according to the rumor, Surditas may cause only minor harms. And as the description of the case makes clear, if the rumor had been that Surditas may cause total deafness, Dr. Negligent would have read the article.

Drug-Induced Deafness shows that an agent may be blameworthy for her benighting act and satisfy the causal condition and the basing constraint (as well as the wrongness condition) without being blameworthy for her unwitting act. However, the case does not show that the revised account is inadequate. Robichaud and Wieland's discussion ignores the fact that the epistemic condition is not satisfied in this case: Dr. Negligent does not know that her not reading the article risks causing deafness in her patients. Hence, rather than showing the inadequacy of the revised account, *Drug-Induced Deafness* highlights the importance of the epistemic condition.

It is worth noting that Robichaud and Wieland do consider a similar epistemic condition earlier in their essay. But they reject that condition on the grounds that in a case analogous to *Extended Stay*, the condition entails indirect blameworthiness where intuitively there is none. (Recall that in *Extended Stay*, although Slacker Dad knew that staying at the park for an extended period of time could lead to a fight between the girls, he is not blameworthy for that fight.) But Robichaud and Wieland make the same mistake here as they do regarding *Drug-Induced Deafness*: like the basing constraint, the epistemic condition is merely a *necessary* condition for indirect blameworthiness. There is no reason why these two necessary conditions (along with other necessary conditions) cannot be combined in a complete account of the *sufficient* conditions for indirect blameworthiness. In some cases (e.g. *Drug-Induced Deafness*), the agent is

not indirectly blameworthy because the epistemic condition is not satisfied; in others (e.g. *Extended Stay*), the agent is not indirectly blameworthy because the basing constraint is not satisfied.

What is Robichaud and Wieland's positive account? They hold what they call a *concern constraint*. According to this constraint, an agent is indirectly blameworthy for an unwitting act only if her benighting act expresses a deficit of concern for the same consideration in virtue of which the unwitting act is wrong. The concern constraint, which is a kind of quality of will condition, requires the agent's benighting act to manifest insufficient concern for the very feature that makes the unwitting act wrong. The constraint yields the right results in *Wrong Prescription* and *Drug-Induced Deafness*. Dr. Lazy's failure to read the article about *Inscientium* manifests a deficit of concern for her patients' well-being, and prescribing *Inscientium* is wrong because it threatens her patient's well-being. Such a match does not obtain in *Drug-Induced Deafness*: Dr. Negligent's failure to read the article about *Surditas* manifests a deficit of concern for mild hearing loss in her patients; however, prescribing *Surditas* to her patient is wrong because it causes total deafness. At this point, I have two remarks to make about Robichaud and Wieland's proposal. First, the cases they present do not warrant the adoption of the concern constraint. I have shown how the revised account respects intuitions about their cases just as well as the concern constraint does. My second point is that although a quality of will condition is not required to account for derivative culpability, it should not be excluded from such an account. In the next section, I will show how a quality of will condition may be integrated (as a potential alternative to an epistemic condition) into an account of derivative culpability. However, in Section 8, I will return to the concern constraint, Robichaud and Wieland's particular version of the quality of will condition, and explain why it is inadequate.

Let us now move to the second concern about the revised account. One may worry that the epistemic condition of the revised account is too cognitively demanding. But this is not the case. Knowing that one's act has a wrong-contributing risk feature need not require a very specific attitude about the nature of that risk. In *Wrong Prescription*, for example, all Dr. Lazy may know is that her failure to read the article about *Inscientium* could lead to a less than adequate treatment of her patients. Her culpability for her omission may not require an attitude with a more specific content. In *Drug-Induced Deafness*, the content of Dr. Negligent's attitude is more specific, though, since the rumor concerns temporary and mild hearing loss. Hence, she knows that her not reading the article risks causing minor harms in her patients, but does not know that it risks causing deafness.

These remarks about the content of the agent's attitude are relevant to a recent debate concerning tracing. To raise doubt about tracing, Manuel Vargas (2005) presents the case of *Jeff the Jerk*. Jeff, a middle manager in a company, is asked to inform a group of employees that they are being laid off. Jeff is a jerk,

and conveys his message to them in a rude and insensitive fashion. Now, Jeff acquired his jerky character when he was a teenager. His goal then was to attract the attention of the opposite sex, and his strategy turned out to be successful.

Intuitively, Vargas notes, Jeff is to blame for laying off his employees in the way he did. However, he adds, this is not a case of direct culpability: due to the specific nature of his jerky character, Vargas stipulates, Jeff cannot help acting the way he does. His blameworthiness must thus be traced to his earlier efforts at becoming a jerk. But Vargas doubts that this tracing condition is satisfied, for during his adolescence, Jeff could not have reasonably foreseen that his acquiring a jerky character would lead to treating his employees in the precise manner in which he did on that day.

In response to Vargas's case, John M. Fischer and Neal Tognazzini (2009) plausibly remark that all that tracing requires in this case is that Jeff could have reasonably foreseen that his adopting a jerky character might result in his treating some people poorly at some point in the future. Just like Dr. Lazy in *Wrong Prescription*, at the time of his benighting act, Jeff need not have a specific attitude about the nature of the future risk.

7. The culpability condition

In the previous section, I remarked that the revised account needs some tweaking. This is because the culpability condition is too weak.

Poisonous Punch. Secret agent Inept is told by her commander that Kurtz will try to kill the president this evening at the end of the Annual Gala. She is instructed to thwart Kurtz's plan. Inept works undercover as a server at the pre-gala reception that Kurtz is expected to attend. Kurtz, Inept knows, loves punch. The ideal course of action for her would be to spike one glass of punch with nausea-inducing bacteria and offer it to Kurtz. Drinking that glass would incapacitate Kurtz and prevent him from doing any harm. But instead, Inept ineptly decides to lace the whole bowl of punch with the bacteria. Kurtz drinks a glass of punch and soon develops incapacitating nausea. Unfortunately, many other guests are similarly indisposed. Now, it turns out that Inept's commander lied to her in order to sabotage her career: Kurtz never planned to kill the president.

Contrary to what the revised account entails, Inept is not derivatively to blame for Kurtz's nausea. Let us examine each condition, starting with the causal one. Clearly, Inept's poisoning of the punch gave rise to Kurtz's nausea. Consider next the wrongness condition. It is wrong for Inept to lace the bowl of punch, because it risked causing nausea in many innocent guests (in the way it did). Moreover, potentially making Kurtz nauseated contributes to making Inept's act wrong.²² Finally, the culpability condition is satisfied. First, Inept is blameworthy for spiking the bowl of punch, for she needlessly exposed innocent guests to the poisonous punch. And second, she knew that spiking the punch risked giving rise to Kurtz's nausea (in the way it did).

Inept is not derivatively to blame for Kurtz's nausea. If he were really planning to kill the president, making Kurtz nauseated would have been the right thing to do. (I am assuming that alternatives such as apprehending Kurtz or physically restraining him are out of the question.) Moreover, Inept had good reasons to believe that Kurtz was indeed planning to assassinate the president. Inept has met her procedural epistemic obligations in this case. We may for instance assume that she has carefully read all the available records concerning Kurtz's plan, and reasonably concluded that they strongly support her commander's claims. In other words, her commander's deception has been well staged. This means that Inept is blameless for her ignorance about Kurtz. (Of course, this does not negate the fact that she is blameworthy for lacing the bowl of punch, since that act unnecessarily risked harming other guests.)

To address the problem raised by this case, we need to tie the agent's culpability more closely to an awareness that risking the undesirable outcome contributes to the wrongness of her act. This is what is missing in Inept's case: she did not know that it is wrong to potentially make Kurtz sick. The culpability condition should be more specific: it must be the case that the agent knew that her act had a wrong-contributing feature consisting in the fact that it risked giving rise to the undesirable outcome in the way it actually did.

Now, I promised at the beginning of this paper that my account of indirect culpability would be neutral on the conditions for direct culpability. I broke this promise in the previous paragraph. As a matter of fact, I broke it as soon as I included Smith's epistemic requirement on direct culpability in the conditions for indirect culpability. Now, one might argue that *knowledge* of wrongdoing – or that one's act has a certain wrong-contributing feature – might not be necessary for direct culpability. Plausibly, an agent who truly believes, but does not know – because of a lack of justification – that his action is wrong or has a wrong-contributing feature is directly blameworthy for performing it.²³ Arguably, then, 'truly believed' should be substituted for 'knew' in the epistemic requirement. Some might endorse an even weaker condition: to avoid blame, they would argue, one must believe that one's action is morally permissible. On this view, an agent would be indirectly culpable for an undesirable outcome only if while performing the benighting act the agent lacks the belief that it has no wrong-contributing risk features. None of these epistemic requirements are fulfilled by Inept, since, by assumption, she reasonably believes that potentially making Kurtz nauseated is a good thing.

It is also worth noting that on some accounts, direct blameworthiness involves *no* epistemic requirement. Quality of will approaches deem an agent blameworthy for her wrongdoing just in case she acts out of ill will or from a lack of good will. According to a version of this view defended by Nomy Arpaly and Timothy Schroeder (2014, Chap. 7), to act out of ill will is (roughly) to act from a desire to bring about wrong-making features, and to act from a lack of good will is (roughly) to act from an insufficiently strong desire to promote good-making

features.²⁴ No epistemic requirement is needed over and above these conditions. On this quality of will view, Inept's poisoning of the punch manifests a lack of good will toward the guests in general, but not toward Kurtz. Since Inept has good reasons to believe that Kurtz is about to do something much worse, her desire to poison him does not constitute ill will or a lack of good will toward him.

To keep my initial promise, my account of indirect culpability should be compatible with all of these perspectives on direct culpability. To maintain neutrality, I will state the culpability condition in terms of the locution 'responds inappropriately to a wrong-contributing feature,' which is meant to allow for various interpretations. To respond inappropriately to a wrong-contributing feature may amount to acting while knowing (or truly believing) that one's act has the wrong-contributing feature. It may also consist in acting while lacking the belief that one's act does not have the wrong-contributing feature, or in acting based on a desire to bring about the wrong-contributing feature or an insufficiently strong desire to avert it, or something else.

I have explained why Inept did not respond inappropriately to the wrong-contributing of her act consisting in the fact that it risked making Kurtz nauseated in the way it actually did. It is worth examining a couple of cases in which agents do respond inappropriately. In *Wrong Prescription*, the crucial wrong-contributing feature of Dr. Lazy's failure to read the article about Inscientium is that this risks harming her patients. She responds inappropriately to this feature, since she is aware of this risk. Alternatively, a quality of will view would deem her response inappropriate because it manifests a lack of good will toward her patients. In *Forest Fire*, the wrong-contributing feature of Hiker's not putting out the campfire consists in the fact that it risked causing a forest fire in the way it actually did. Hiker responds inappropriately to this risk, since he knows about it. Proponents of a quality of will approach would contend that his response is inappropriate in virtue of the fact that it manifests a lack of good will toward those who would be negatively affected by the fire.

We are now in a position to formulate a revised culpability condition that is neutral on the conditions for direct culpability: an agent is derivatively culpable for an undesirable outcome only if she is directly culpable for the benighting act, and while performing that act, she responded inappropriately to its wrong-contributing feature consisting in the fact that it risked giving rise to the undesirable outcome in the way it actually did.

8. Simplifying the account

I have argued that an agent is indirectly to blame for an undesirable outcome just in case (1) the agent performed a benighting act that is wrong at least in part because it risked giving rise to the undesirable outcome in the way it actually did, (2) the benighting act gave rise to the undesirable outcome, and (3) the agent is directly culpable for the benighting act, and while performing

that act, she responded inappropriately to its wrong-contributing feature consisting in the fact that it risked giving rise to the undesirable outcome in the way it actually did.

Now, let us try to simplify the account a little. The second conjunct of the new culpability condition (3) makes the wrongness condition (1) redundant, on the understanding that the locution 'responds inappropriately' is *factive*, that is, the truth of 'S responds inappropriately to x' entails the occurrence of x.²⁵ (I should also remind the reader that according to a stipulation I made at the end of Section 4, any act that has a wrong-making or wrong-contributing feature is automatically morally wrong.) Indirect culpability still requires the satisfaction of the wrongness condition, but we can now take this condition to be integrated into the culpability condition. We thus have: an agent S is indirectly to blame for an undesirable outcome O just in case (1) S's action A caused O, and (2) S is directly culpable for A, and while performing A, S responded inappropriately to A's wrong-contributing feature consisting in the fact that it risked causing O in the way it actually did.

It is useful to compare this account to Robichaud and Wieland's. Their concern constraint, let us recall, holds that an agent is indirectly blameworthy for an unwitting act only if her benighting act expresses a deficit of concern for the same consideration in virtue of which the unwitting act is wrong. The concern constraint is similar to (2), but there are several important differences. First, as I argued in Section 6, the cases Robichaud and Wieland present do not warrant the adoption of the concern constraint. More generally, the question whether direct culpability requires a quality of will condition is contentious, and an account of indirect culpability should steer clear of this debate. For this reason, the neutral language of (2) is preferable. (2) is more general than the concern constraint in another respect, since it leaves open the specific nature of the undesirable outcome O: O may be an unwitting act, but it may also be an incompetent act or a harmful event. I will return to the concern constraint shortly, and point out two additional issues with it.

At the beginning of this essay, I wrote that the three conditions on derivative culpability ought to be more specific and exhibit a more intimate relationship with each other than is usually thought. (2) holds a kind of *matching constraint* concerning the three components of indirect culpability: S's wrongdoing A, S's culpability for A and A's causing the undesirable outcome O. For indirect culpability to obtain, the feature that (in part) makes A wrong must match the feature that (in part) makes S blameworthy for A, and that feature must concern the risk of causing O in the way it actually causes it. The wrongness condition does not merely require the moral wrongness of action A: it requires A to be wrong at least in part because A risked causing O in the way it actually did. Slacker Dad (of *Extended Stay*) does not satisfy this condition: what makes his decision to stay at the park wrong is the fact that it breaks his promise to his wife rather than the fact that it could lead to a fight between the girls. Dr. Lazy (of *Wrong Prescription*)

does satisfy the condition: her failure to read the article about Inscientium is wrong because it risked harming her patients. The condition is also specific about what makes *S* blameworthy for *A*: *S*'s blameworthiness must be due (at least in part) to *S*'s inappropriate response to the fact that *A* risked causing *O* in the way it actually did. Prof. Gamer (of *Blood Type*) does not satisfy this condition: his failing to survey the members of his department about their blood types is not an inappropriate response to his colleagues' well-being, but rather an inappropriate response to the promise he made to his son. Drunk Driver (of *Drink and Drive*) does satisfy the condition: his getting drunk is an inappropriate response to the potential danger his driving while impaired will cause to others.

The concern constraint is also a matching constraint: according to this constraint, the feature that the agent's benighting act manifests insufficient concern for ought to match the feature that makes the unwitting act wrong. But this constraint is too rough. *Multiple Side Effects* shows that instead of invoking wrong-making features, the condition on indirect culpability should invoke wrong-contributing features. Second, and more importantly, the concern constraint does not offer an adequate treatment of cases such as *Questionable Lesson*. These cases illustrate the fact that there are potential trade-offs between the wrongness condition and the culpability condition. According to most normative theories, Hiker's omission is not wrong in virtue of the fact that it risked causing the forest fire in the way it actually did. However, according to actual consequentialism, it is. However, on the latter view, Hiker is not indirectly culpable for the forest fire, because he did not respond inappropriately to the fact that a forest fire risked being caused in that highly unexpected way. Now, the concern constraint applies only to instance of indirect culpability involving culpable ignorance. But we can easily modify the constraint to make it applicable to other instances of indirect culpability: an agent is indirectly blameworthy for an undesirable outcome only if her benighting act expresses a deficit of concern for what makes the outcome undesirable. Hiker, in *Questionable Lesson*, satisfies this condition, regardless of how objective one's normative theory is. This means that Hiker would be blameworthy for the forest fire, if we accept both actual consequentialism and this revised concern constraint. Note that requiring that the causal path between the benighting act and the undesirable outcome be non-deviant does not help here, since, as we saw, actual consequentialists consider that path to be non-deviant.

It is worth noting that this kind of issue arises also with respect to non-radical objective ethical theories. Consider a case like *Forest Fire* and *Questionable Lesson*, except that the causal path leading to the forest fire is rather unusual, but not completely fluky. More specifically, the path is expectable, given a certain body of evidence, but not expectable, given Hiker's evidence. Suppose further that Hiker's ignorance is not itself culpable. Hence, intuitively, Hiker is not blameworthy for the forest fire in this variant of the case. His failure to douse the campfire is wrong and culpable because there is a number of ways in which,

Hiker knew, it could have led to a forest fire. Is the causal path that actually obtained one of the ways that make his omission wrong? In other words, is the actual causal path non-deviant? According to some normative theories, the answer is 'no.' This verdict would explain why Hiker is not blameworthy for the forest fire: it is not the case that Hiker's omission is wrong because it risked giving rise to the forest fire in the way it actually did. But according to more objective theories, the answer is 'yes.' To explain why Hiker is not blameworthy for the forest fire, these theories need to invoke a culpability condition like (2). Hiker is not to blame for the forest fire, because it is not the case that he responded inappropriately to the fact that his omission risked causing a forest fire in the way it actually did. Now, according to these more objective theories, both a causal condition requiring a non-deviant path and the concern constraint would be satisfied in this case. Proponents of the concern constraint would thus have to deem Hiker blameworthy for the forest fire, contrary to intuition.

9. Conclusion

Here is, one last time, my account: an agent *S* is indirectly to blame for an undesirable outcome *O* just in case (1) *S*'s action *A* caused *O*, and (2) *S* is directly culpable for *A*, and while performing *A*, *S* responded inappropriately to *A*'s wrong-contributing feature consisting in the fact that it risked causing *O* in the way it actually did. I have shown how this account respects intuitions about a wide variety of cases. Crucially, the account steers clear of controversies in two related but distinct areas. It is neutral about normative ethical issues: it is the role of a normative theory to determine whether the fact that *A* risked causing *O* in the way it actually did is wrong-contributing or not. Second, the locution 'responds inappropriately' is meant to be neutral about the specific motivations from which *S* performed *A*.

Notes

1. See, among others, Vargas (2005), Fischer and Tognazzini (2009), Timpe (2011), Khoury (2012) and Shabo (2015).
2. On some versions of the story, his culpability is traced to his decision to drive the car while drunk. But I am assuming that not driving his car home is not an option.
3. See, among others, Strawson (1962), Wallace (1994, Chap. 3) and McKenna (2012, Chap. 2). Even consequentialists can coherently embrace this distinction between culpability and moral sanction. As Nomy Arpaly (2003, 70) usefully reminds us, consequentialists need not analyze every moral assessment in terms of consequences. Their position about the moral wrongness of an action is compatible with holding that blameworthiness should be understood in terms of the agent's motives rather than, say, the consequences of our overtly blaming her.
4. See (Smith 1983, 564).
5. However, Adams (1985) and Smith (2005) argue that one may be directly blameworthy for one's emotions, beliefs and other attitudes. Although I disagree

with this position, I grant that the issue is complex and would deserve more scrutiny than I can offer here.

6. See also Feinberg (1970) and Rosen (2008, 609–610). Fischer and Ravizza (1998, Chap. 4) consider responsibility for consequences. Although they do not explicitly hold that agents are indirectly responsible for the consequences of their actions, their account does not prohibit this interpretation.
7. Some authors hold that an agent may be blameworthy for a morally permissible act. See, for instance, Zimmerman (1997), Haji (1998, 148–150), Vranas (2007, 193–196) and Capes (2012). These authors may be inclined to hold that an agent could also be indirectly blameworthy for the bad consequences resulting from a permissible action. I do not have the space to discuss this position; however, in Montminy (2018), I defend the view that one may not be blameworthy for a morally permissible act.
8. See, among many others, Smith (1983), Rosen (2004, 299), Widerker (2006, 163, n. 1) and Fischer and Tognazzini (2009, 532).
9. Similar cases are presented in Feinberg (1970).
10. See Feinberg (1970) and (Ginet 2000) for similar accounts.
11. See Feinberg (1970, 196) for a similar diagnosis on a similar case.
12. See Feinberg (1970, 196, 199) for similar remarks.
13. See, for instance, (Armstrong 1999); (Beebe 2004).
14. See, for instance, Dowe (2001). For our purposes, it will not matter whether a quasi-cause is a genuine type of cause or not a cause at all.
15. Why 'roughly'? Simple counterfactual accounts of causation (or quasi-causation) need to be refined to deal with cases of overdetermination, preemption and the like.
16. Smith (*ibid.*, 556) briefly considers deviant causal chains, but only with respect *direct* culpability. To be directly to blame for an action, she remarks, the agent's problematic motives must give rise to the action in the right kind of way.
17. Feinberg's (1970) account of derivative culpability involves a sophisticated causal condition. Consideration of cases similar to *Extended Stay* and *Blood Type* leads Feinberg to admit a wrongness condition very much like the one proposed in the previous section. It also leads him to adopt a causal condition according to which the risk of harm that, at least in part, made the benighting act wrong is the risk of just the sort of harm that was in fact caused by the benighting act (*ibid.*, 199). Clearly, both this condition and the wrongness condition are satisfied in *Questionable Lesson*.
18. See Sinnott-Armstrong (2014) for a useful overview.
19. See, among others, (Ross 1939); (Prichard 1968); (Zimmerman 2008).
20. See Smith (2010, 2014) for a useful overview and criticisms.
21. This verdict would be rejected by subjective ethical theories, but let us set that issue aside.
22. This verdict does not require us to adopt actual consequentialism, or some radical objective theory. We can suppose that it is well known that Kurtz is an ordinary, law-abiding citizen who is not involved in any criminal activity; however, Inept's commander and colleagues have pulled an elaborate and persuasive prank that has convinced her that he is a professional assassin.
23. See (Rosen 2008, 596–597) for similar remarks. See also Haji (1997), Ginet (2000), Zimmerman (2008), Sher (2009), Harman (2011), Levy (2011) and (Montminy 2016), and for recent discussions of the epistemic condition.
24. See also Markovits (2010), (Smith (2011) and McKenna (2012)). Robichaud and Wieland's concern constraint, discussed in Section 6, is a version of this approach.

25. The causal condition is also redundant; however, it seems preferable to keep an explicit statement of that condition.

Acknowledgement

I am grateful to anonymous referees for this journal for their valuable comments on earlier drafts of this paper.

Notes on contributors

Martin Montminy is Professor of Philosophy at the University of Oklahoma. He works in the philosophy of language, epistemology and metaphysics. His recent publications have been on topics such as causation, vagueness, moral responsibility and the norms of assertion.

References

- Adams, Robert. 1985. "Involuntary Sins." *Philosophical Review* 94: 1–31.
- Armstrong, David. 1999. "The Open Door." In *Causation and Laws of Nature*, edited by H. Sankey, 175–185. Dordrecht: Kluwer Academic.
- Arpaly, Nomy. 2003. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press.
- Arpaly, Nomy, and Timothy Schroeder. 2014. *In Praise of Desire*. Oxford: Oxford University Press.
- Beebe, Helen. 2004. "Causing and Nothingness." In *Causation and Counterfactuals*, edited by J. Collins, N. Hall, and L. A. Paul, 291–308. Cambridge (MA), The MIT Press.
- Capes, Justin. 2012. "Blameworthiness Without Wrongdoing." *Pacific Philosophical Quarterly* 93: 417–437.
- Dowe, Phil. 2001. "A Counterfactual Theory of Prevention and 'Causation' by Omission." *Australasian Journal of Philosophy* 79: 216–226.
- Feinberg, Joel. 1970. "Sua Culpa." In *Doing and Deserving*, edited by J. Feinberg, 187–221. Princeton (NJ), Princeton University Press.
- Fischer, John Martin, and Mark Ravizza. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fischer, John Martin, and Neal A. Tognazzini. 2009. "The Truth about Tracing." *Noûs* 43: 531–556.
- Ginet, Carl. 2000. "The Epistemic Requirements for Moral Responsibility." *Philosophical Perspectives* 14, *Action and Freedom*: 267–277.
- Haji, Ishtiyaque. 1997. "An Epistemic Dimension of Blameworthiness." *Philosophy and Phenomenological Research* 57: 523–544.
- Haji, Ishtiyaque. 1998. *Moral Appraisability: Puzzles, Proposals and Perplexities*. Oxford: Oxford University Press.
- Harman, Elizabeth. 2011. "Does Moral Ignorance Exculpate?" *Ratio* 24: 443–468.
- Khoury, Andrew. 2012. "Responsibility, Tracing, and Consequences." *Canadian Journal of Philosophy* 42: 187–207.
- Levy, Neil. 2011. *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. Oxford: Oxford University Press.
- Markovits, Julia. 2010. "Acting for the Right Reasons." *Philosophical Review* 119: 201–242.

- McKenna, Michael. 2012. *Conversation and Responsibility*. Oxford: Oxford University Press.
- Montminy, Martin. 2016. "Doing One's Reasonable Best: What Moral Responsibility Requires." *Journal of the American Philosophical Association* 2: 55–73.
- Montminy, Martin. 2018. "Culpability and Irresponsibility." *Criminal Law and Philosophy* 12: 167–181.
- Prichard, H. A. 1968. "Duty and Ignorance of Fact." *Moral Obligation and Duty and Interest*, 18–39. Oxford: Oxford University Press.
- Robichaud, Philip and Jan Willem Wieland. *Forthcoming*. "A Puzzle Concerning Blame Transfer." *Philosophy and Phenomenological Research*.
- 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.
- Ross, W. D. 1939. *Foundations of Ethics*. Oxford: Clarendon Press.
- Shabo, Seth. 2015. "More Trouble with Tracing." *Erkenntnis* 80: 987–1011.
- Sher, George. 2009. *Who Knew? Responsibility without Awareness*. Oxford: Oxford University Press.
- Sinnott-Armstrong, Walter. 2014. "Consequentialism." *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. <http://plato.stanford.edu/archives/spr2014/entries/consequentialism/>
- Smith, Angela. 2005. "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115: 236–271.
- Smith, Holly. 1983. "Culpable Ignorance." *The Philosophical Review* 92: 543–571.
- Smith, Holly. 2010. "Subjective Rightness." *Social Philosophy and Policy* 27: 64–110.
- Smith, Holly. 2011. "Non-Tracing Cases of Culpable Ignorance." *Criminal Law and Philosophy* 5: 115–146.
- Smith, Holly. 2014. "The Subjective Moral Duty to Inform Oneself before Acting." *Ethics* 125: 11–38.
- Strawson, Peter. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48: 1–25.
- Timpe, Kevin. 2011. "Tracing and the Epistemic Condition on Responsibility." *The Modern Schoolman* 88: 5–28.
- Vargas, Manuel. 2005. "The Trouble with Tracing." *Midwest Studies in Philosophy* 29: 269–291.
- Vranas, Peter. 2007. "I Ought, Therefore I Can." *Philosophical Studies* 136: 167–216.
- Wallace, Jay. 1994. *Responsibility and the Moral Sentiments*. Cambridge (MA), Harvard University Press.
- Widerker, David. 2006. "Libertarianism and the Philosophical Significance of Frankfurt Scenarios." *Journal of Philosophy* 103: 163–187.
- Zimmerman, Michael. 1997. "A Plea for Accuses." *American Philosophical Quarterly* 34: 229–243.
- Zimmerman, Michael. 2008. *Living with Uncertainty*. Cambridge: Cambridge University Press.