

## IN DEFENSE OF MODERATION

### CULPABLE IGNORANCE AND THE STRUCTURE OF EXCULPATION

*James Goodrich*

**A**N INFANT begins to drown in a neighborhood pool. Anne—the local on-duty lifeguard—jumps in, pulls the infant out, and performs standard CPR. The infant dies. The infant would have survived had Anne performed a different CPR procedure—the one designed for infants. Why did Anne not perform the correct procedure? She did not know about it. And she did not know about the special procedure for infant CPR because, during her training, Anne left early to take a smoke. Anne knew that important information might be shared during the ten minutes she was gone, but took the risk anyway.<sup>1</sup> Anne was thus knowingly reckless and culpable for her ignorance. Anne’s recklessness, moreover, at least partially explains why she failed to save the infant. Does this mean that Anne is culpable to some degree for failing to save the infant’s life?

Let us clarify this question with a distinction.<sup>2</sup> *Benighting acts* are those acts in which the culpably ignorant agent culpably fails to remedy her ignorance or risks missing out on some morally important information that might help guide her future decisions. Anne’s act of sneaking out for a cigarette was a benighting act. The *unwitting wrongful act* is the later, objectively wrong act that the culpably ignorant agent performs out of their ignorance. Anne’s unwitting wrongful act was her failure to save the infant’s life.

Culpably ignorant agents, by definition, are culpable for their benighting acts. But there is disagreement over whether culpably ignorant agents are also culpable for their unwitting wrongful acts. Liberals think they are not. Conservatives and moderates believe they are culpable to at least some degree for their unwitting wrongful acts. The difference is that while conservatives believe agents are fully culpable for their unwitting wrongful acts, moderates believe

1 This case is inspired by Smith, “Culpable Ignorance,” 552.

2 This terminology is introduced in Smith, “Culpable Ignorance.”

that they are culpable to some degree, though not necessarily fully culpable for their unwitting wrongful acts.

I like the moderate view. It best accords with my—and I suspect others’—intuitions.<sup>3</sup> However, liberal critics have mounted a difficult and hitherto unanswered challenge to the moderate view. Roughly: the moderate must explain why an agent can be culpable (to some degree) for their unwitting wrongful acts because they are culpable for their benighting acts. And this explanation must sit well with plausible accounts of culpability. It has proven more difficult than one might have thought to meet this challenge.

I will defend the moderate view against the liberal’s challenge. I will begin by developing a novel account of three things: (1) the grounds of culpability, (2) the grounds of excuses, and (3) the way excuses function within a theory of culpability. On my view, culpability is grounded in facts about wrongdoing *tout court*. However, the culpability-grounding function of facts about wrongdoing can be disabled by undercutting defeaters. These undercutting defeaters are what we colloquially refer to as “excuses” and excuses are then grounded in facts about an agent’s quality of will. If I am right, the liberal’s challenge hinges upon unchecked philosophical assumptions about the nature and structure of culpability.<sup>4</sup>

In the first four sections, I clarify the nature of the problem that animates my search for a new theory of the relationship between culpability and excuse and outline my account. Sections 5 through 8 develop my new account and defend it against objections.

### 1. THE LIBERAL’S CHALLENGE

Here is the liberal’s challenge: moderates must explain why culpably ignorant agents are only partially excused for their unwitting wrongful acts in light of standard accounts of culpability.<sup>5</sup> In particular, the moderate needs to explain

3 There is room for disagreement between moderates about particular cases, including the one I have presented here. I propose that we grant that the moderate believes Anne to be culpable to some degree for her unwitting wrongful act. After all, this case is relevantly like the cases Smith deploys. See, again, Smith, “Culpable Ignorance,” 556. If the moderate can meet the liberal’s challenge on the very set of cases for which the liberal thinks their challenge best applies, we will ensure that no questions are begged.

4 Throughout this essay, I will only be concerned with ignorance of descriptive facts.

5 See especially Smith, “Culpable Ignorance,” and “Tracing Cases of Culpable Ignorance”; and Husak, *Ignorance of Law*, ch. 3.

how their view is compatible with a quality of will account of culpability (QWA).<sup>6</sup> Holly Smith puts the point as follows:

The culpably ignorant agent cannot be held to blame for his unwitting act, since he fails one of the conditions of culpability. His act does not arise from a defective configuration of desires and aversions.<sup>7</sup>

The thought is this: for an agent to be culpable, the wrong action for which the agent is culpable needs to have been produced by a morally objectionable motive, intention, or desire. And this is meant to fall out of our best account of what it is for an agent to be culpable. Smith characterizes her version of this account (roughly) as follows:

*Smith's QWA:* The fact <S is culpable for A> is grounded in the facts that

1. <A (or its attempt) is objectively wrong>,
2. <S had a reprehensible configuration of desires and aversions>,  
and
3. <This configuration gave rise to the performance of A>.<sup>8</sup>

To get a feel for why one might be attracted to this view, consider a toy case: Beth is a conscientious walker who trips on an uneven sidewalk, falls into a puddle, and thereby splashes muddy water onto a passerby. It seems inappropriate to blame Beth. She might apologize out of kindness, but would surely be right to say, "I didn't mean to!" After all, her action did not arise out of some motivation or intention to harm or do wrong to the passerby. In fact, it did not even seem like an action! It was just an accident. Indeed, the appeal to "I didn't mean to!" is an expression of the fact that the agent did not intend any harm (and perhaps that she therefore should not be blamed).

Reconsider Anne, the lifeguard. When Anne performs adult CPR on the infant, what does she intend to do? In the version of the story I have offered, it seems like she intends to save the child. This intention is good, even noble. The fact that Anne fails to save the child (or even hastens its death) is antithetical to her intended aims. In other words still, Anne's failure to save the infant did not arise out of a reprehensible configuration of desires. Her unwitting wrongful

6 For a survey of recent quality of will accounts, see Shoemaker, "Qualities of Will."

7 See Smith, "Culpable Ignorance," 559. For similar remarks see Smith, "Tracing Cases of Culpable Ignorance," 113.

8 See Smith, "Culpable Ignorance," 556. Note that Smith does not state her account in terms of grounding conditions, but in terms of truth conditions. This is plausibly due to the philosophical norms of the era. In her writing, Smith is clearly concerned with something like explanation, not with what makes a sentence true or false. Smith confirms this in correspondence.

action arose out of good motives. If it is a necessary condition on an agent being culpable that their wrong act arises from some objectionable motives, culpably ignorant agents are not culpable at all for their unwitting wrongful acts. And this contradicts the moderate's view.

The moderate could respond by rejecting QWAs of culpability *tout court*. Some may be attracted to this move. I am not. I would rather mount a defense of the moderate view that does not crucially turn on whether we should accept a QWA of culpability. To do so would be to tether the fate of the moderate view to the hope that no form of the QWA will win the battle of theoretical virtues. Moreover, for such a move to help the moderate, it would also need to be shown that similar challenges do not arise on other accounts of culpability. As someone who is a moderate first, I would rather not take that gamble.

Here is a different move the moderate could make: the moderate could appeal to the objectionable motives that gave rise to culpably ignorant agents' benighting acts. These earlier motives are clearly objectionable. Anne should not, for example, go out for a smoke during lifeguard training. Perhaps the moderate can then say that an agent's unwitting wrongful act did arise out of a morally objectionable motive in the following sense:

*Transfer Model:* Morally objectionable motives can “morally transfer” across (the right sort of) causal relations to give rise to later actions.

If the moderate does adopt the Transfer Model, perhaps they can explain why culpably ignorant agents are at least partially culpable for their unwitting wrongful acts: their earlier morally objectionable motives are causally related to their unwitting wrongful act.

The liberal could object that the Transfer Model leaves everything to be explained. What theory-neutral reasons do we have to think that such causally distant motives are to be considered in determining the culpability of a given action? From the liberal's point of view, the Transfer Model may seem like little more than a restatement of the moderate's intuition suitably dressed for a QWA of culpability. What the moderate needs is a credible, theory-neutral rationale for something like the Transfer Model. Without such a rationale, the moderate is open to the criticism that the Transfer Model merely reasserts the intuition that the liberal is inclined to reject. Moreover, plausibly not all motives in the causal chain leading up to a particular action are relevant. The “right sort of” locution invites reasonable philosophical suspicion. Why think that there is a way to characterize the “right sort of” causal relations such that they amount to more than the causal relations that fit the moderate's intuitions?

The Transfer Model therefore does not explain everything that needs explaining. Why would causally or temporally distant motives count as “giving

rise to” actions in the morally relevant sense? And why would some causally or temporally distant motives count as giving rise to an action, but other motives would not? The moderate needs more than an intuition here. The liberal and moderate start from a clash of intuitions. The liberal seems to be winning by incurring fewer explanatory peculiarities. The moderate needs a response to these explanatory challenges that does not commit them to such peculiarities or to claims the liberal could insist we not make.

## 2. MY STRATEGY FOR DEFENDING THE MODERATE VIEW

The liberal’s challenge to the moderate view is that the moderate’s explanation of why the culpably ignorant agent is culpable looks unmotivated when combined with a QWA of culpability. How should moderates respond?

I will accept the liberal’s presupposition about cases like that of Anne the lifeguard. Namely, there is not a bad motive that plausibly gives rise to Anne’s unwitting wrongful action. I think we can reply to the liberal’s challenge while holding this assumption fixed.

However, I will argue that moderates can help themselves to an alternative version of the QWA for which the liberal’s explanatory challenge does not arise. That is, the version of the QWA that leads to the liberal’s challenge has some unchecked theoretical baggage. Once we see that this theoretical baggage is unnecessary, we will see how the moderate can answer the liberal’s challenge. Therefore, there is a version of the QWA of culpability that, by the moderate’s own lights, runs into no serious explanatory challenges. However, I will not argue that the liberal must accept my new version of the QWA. Rather, my point is that the liberal’s challenge relies on an inference from the claim that the moderate view faces explanatory trouble on one plausible QWA to the claim that the moderate view faces explanatory trouble on all plausible QWAs. My point then is that this inference at the heart of the liberal’s challenge is unwarranted.

My account of culpability, *contra* Smith’s, does not require an appeal to the “gives rise to” relation. The liberal’s challenge is about how to make sense of objectionable motives giving rise to objectionable actions when such motives are causally distant or perhaps not even plausibly part of the causal chain. If, however, we can have a theory that does not rely on this notion as necessary, then we incur no explanatory burden. That is, we should do without the Transfer Model because it gives the “gives rise to” relation a central explanatory role. To be clear: I leave it open whether the “gives rise to” relation does work in some cases. My point is only that it is not necessary.

However, I cannot simply stipulate that I am dispensing with the claim that the “gives rise to” relation is necessary for grounding culpability. If I could,

this essay would end here. Why? Because the “gives rise to” relation has great explanatory power and it is not obvious what a satisfying account of culpability that dispenses with its central importance looks like. I will therefore need to replace the centrally important “gives rise to” relation with some other, plausible-enough theoretical tools. And these tools had better not just be cognates for which the liberal’s challenge arises all over again. Thus, to credibly dispense with the “gives rise to” relation, I must answer some more foundational questions. These questions are about the essential nature of culpability and its grounds. I will now turn to constructing this theory. It will be best to proceed in small steps, for the QWA often assumed by liberals is, in some ways, only subtly different than my own. But these subtle differences, taken in conjunction, make all the difference to the plausibility of the moderate view.

### 3. THE SPARSE THEORY

My alternative proposal relies on a division of theoretical labor. I will first offer a theory of what grounds the fact that a given agent is culpable for a given action; then I will offer a theory of how excuses can “swoop in” to get agents off the culpability hook. Being clear about the difference between the grounds of culpability and the explanatory structure of excuses is key in understanding how we can avoid invoking the “gives rise to” relation.

Here is my account of the grounds of culpability:

*The Sparse Theory:* The fact <S is culpable for A> is sometimes explained by the fact that <A (or its attempt) is morally wrong>.

That is it. The primary explanatory fact in my account of culpability is only that an agent has performed a morally wrong action.<sup>9</sup>

What do I mean by “wrong”? I mean whatever sense of that term you believe is most important in normative ethics. Some place great weight on the “objective” or “fact-relative” sense of wrong. Others will hold that it is more important that we focus on “subjective,” “belief-relative,” or “evidence-relative” conceptions. Still others will be happy to adopt a kind of pluralism, accepting each concept of moral wrongness as equally important and perhaps accepting a distinct concept of culpability corresponding to each sense of wrongness. It matters little for the purposes of this essay which sense of wrongness we deploy.<sup>10</sup>

9 There is a variant of the Sparse Theory according to which it is not the fact that an action is morally wrong that grounds culpability, but rather, it is the facts that make an act wrong that ground culpability. For present purposes, we can be agnostic about which is superior.

10 For what it is worth, I prefer the third, pluralist understanding of wrongness and culpability. The pluralist understanding allows us to describe situations with maximal

Of course, the Sparse Theory would be quite implausible if it were the only thing we said about culpability. After all, were you to perform a wrongful act out of non-culpable ignorance, you would not be culpable. A theory of excuse is needed to explain this fact. The Sparse Theory is therefore far more plausible when supplemented with a theory of excuse. Excuses, on my view, are undercutting defeaters. I will say more about undercutting defeaters in due course, but it is worth first discussing the relationship between the idea that excuses are defeaters and the idea that wrongful actions are the grounds of culpability.

While the fact that some agent performed a wrong action is the grounds of culpability—that is, it is the operative explanatory factor of why a given agent is culpable for a given action—background conditions still need to be met in order for wrongful action to play its grounding role. Consider a common analogy: a match is lit because I struck it. The fact that I struck the match is a perfectly good explanation of the further fact that the match is lit. However, certain background conditions must be in place in order for the striking of the match to successfully light the match: there must be sufficient oxygen in the room, the match must be sufficiently dry, and so on. Though these additional background factors play some role in the fullest possible explanation of the match being lit, such conditions are not the operative ground in question. They instead do something to explain why the operative ground itself was indeed operative in the given context.

This may seem strange. Why should a QWA theorist find the claim that bad motives merely function as explanatory background conditions plausible enough for the purposes of this discussion? Is there not some sense in which the whole point of the QWA is to say that the quality of the agent's will is more like the striking of the match than like the oxygen in the room? However, keep in mind that, according to the Sparse Theory, it is possible that facts about an agent's quality of will do play an operative role in some token explanations of why an agent is culpable. For example, it could well be true that the malicious intent of a murderer plays a role in explaining why they are culpable for murdering someone. My point is rather that this need not be the only way for facts about an agent's quality of will to play a part in a complete explanation for why they are culpable. We can separate out a theory of excuse that involves plausible claims about how facts about an agent's quality of will can figure in as background conditions in the explanation of an agent's culpability.

Thus, one possibility is that the operative grounds of the murder's culpability are overdetermined. Both the wrong action itself (without a further quality

---

specificity. A given agent may act wrongfully in one sense and thus be culpable in the corresponding sense without it being true that they are culpable in any other important sense.

of will excuse to be discussed in the next section) and the malicious motives would be sufficient to ground the murderer's culpability. Insofar as that is true, we can also accept, as some would like to, that agents can be culpable without being culpable for a wrongful action.<sup>11</sup> Those are just cases when only the bad motives are an operative ground of culpability.<sup>12</sup>

#### 4. THE SPARSE THEORY AND EXCUSES

Now, reconsider excuses. We can think of excuses and their role as defeaters in a similar way to the background conditions needed for the match to be lit. Excuses may, in a given context, do something to explain why the fact that a given action is morally wrong fails to ground the further fact that an agent is culpable. But again, we will flesh out this idea in greater detail in what follows. Let us now turn to the question of how thinking about the division of labor between the Sparse Theory and the idea that excuses are defeaters fits in with liberal's challenge to the moderate view.

The Sparse Theory is compatible with any number of different substantive theories of excuse. For our purposes, we will be interested in a quality of will (QOW) theory of excuse. That is, the class of facts that count as excuses, according to my QOW theorist, is characterized by facts about whether the agent under consideration had a reprehensible configuration of desires, intentions, aversions, etc. Though it should be obvious by now, this locates the explanatory importance of the quality of an agent's will in the excuse part of our theory, not in the fundamental grounds of culpability itself. This, I think, is good enough for my view to count as a kind of QWA. But it does differ in structure from what I have called Smith's QWA. I offer a QWA of excuses, not a QWA of the grounds of culpability. Smith offers a QWA of the grounds of culpability.

But how does the Sparse Theory, when combined with an adequate theory of excuse, help the moderate? Recall that, earlier, I said we should get rid of the central explanatory importance of the "gives rise to" relation. The Sparse Theory clearly does that. But how are we going to characterize the class of facts that count as excuses without helping ourselves to the "gives rise to"

11 For discussion, see, e.g., Capes, "Blameworthiness without Wrongdoing."

12 For what it is worth, considerations of parsimony tempt me to the view that wrongful action is always the only operative ground of culpability and facts about an agent's quality of will only come in on the excuse part of the theory. I think Capes's arguments can be handled by appeal to the right theory of subjective wrongness ("Blameworthiness without Wrongdoing"). However, giving that much weight to parsimony and responding to Capes in this way will both be controversial moves among QWA theorists, so they need not follow me on these further controversial claims. They are not central to the topic at hand.



relation? We are not. Here, we will appeal to the “gives rise to” relation. For some fact to excuse some agent for their wrongful action—and thereby prevent the successful grounding of culpability—that fact needs to appeal to some morally acceptable motivations that gave rise to the wrongful action under assessment.

One might ask: Why does my appeal to the “gives rise to” relation in my theory of excuse not just reintroduce the liberal’s challenge under a different guise? There are two parts to my answer: (1) we are using facts about the quality of an agent’s will not to explain why some agent is culpable, but to explain why they are excused; and (2) we can assess the moral acceptability of a given excuse without appealing all and only to the wrongful action. Let us take each of these points in turn.

On the first point: according to Smith, we use facts about the moral acceptability of an agent’s motivations to explain why some agent is culpable. On my view, we use facts about the moral acceptability of an agent’s motivations to explain why some agent is excused. This means that when I appeal to the “give rise to” relation, I am not appealing to it in order to explain why an agent is culpable. So, if I think that culpably ignorant agents are culpable to some degree for their unwitting wrongful action, the ground of their culpability still stems from the wrongness of their action. It does not stem from some causally distant motives. It is therefore not incumbent upon me to explain why causally distant motives would ground the culpability of culpably ignorant agents. When I appeal to the “gives rise to” relation, I am doing so to characterize the class of facts that count as excuses. I therefore have not incurred the sort of explanatory burdens pointed to by the liberal. And this is in part because my explanation of why an agent is culpable does not appeal to the sorts of facts it which the liberal assumes the moderate is appealing.

On the second point: I have divided out the theory of excuse from the theory of what grounds culpability. When an agent does wrong, they are culpable unless some fact—an excuse—defeats the grounding relation given by the wrongful action. Excuses are facts about a morally acceptable set of motivations that gave rise to the wrongful action in question. If someone has an acceptable and morally sincere set of motivations, then they have an excuse for having done wrong. The question that divides the liberal and moderate then becomes “What constitutes a morally acceptable set of motivations?” Part of the answer the liberal and moderate will agree on: the content of the motivations. Desires to harm undeserving people are morally unacceptable. Desires to save people are, *ceteris paribus*, morally acceptable.

But the moderate and liberal disagree about what conditions can affect the moral acceptability of an excuse outside of the content of the motivation.

Moderates think the relations that an otherwise good motive stand in to other mental states may be relevant to whether an agent's otherwise good motive is exculpating. That is, the fact that someone has an otherwise good motive counts as an excuse only if that motive fails to stand in the appropriate relations to other morally objectionable motives. We might think of these other facts about the relations otherwise good motives stand in to other bad motives as defeaters for the facts that would otherwise count as an excuse. In other words still, we can help ourselves to the idea that there is a recursive structure to defeat that tells us something about excuses.

To illustrate this rather abstract point, reconsider Anne. Anne's failure to save the child is a *pro tanto* ground of culpability. This *pro tanto* ground could potentially be defeated by the fact that Anne had a sincere motive to save the child—an excuse. But this *pro tanto* excuse is itself (partially) defeated by the further fact that Anne's motivation to save the child is combined with a belief that had been given rise to by past bad motivations. The important thing to note is that Anne's culpable motivations for her benighting act do not explain why she is culpable for her unwitting wrongful act. They explain why her otherwise noble motives are insufficient as an excuse for wrongdoing.

We can summarize this view (roughly) as follows:

*QOW Excuses:* Some fact  $F$  counts as an excuse for some wrongdoing (and thereby defeats the grounding relation) only if:

1.  $F$  is about a motive that both give rise to  $\phi$  and has morally acceptable content; and
2. There is no further fact  $G$  that defeats the exculpatory force of  $F$ .

It is important for my account that 1 and 2 are clearly separated out, for  $G$  need not itself stand in any special relation to over and above the extent to which  $G$  stands in the right kind of relation to  $F$ . This means that the fact that the culpably ignorant agent has a bad motive at a prior point in time needs to be suitably related to the *pro tanto* excusing fact. It does not itself need to explain anything about why the agent is culpable in the first place. The moderate therefore fails to incur the explanatory challenge they would incur if they were to instead accept Smith's QWA. We will discuss this point in greater substantive detail later in this essay.

Even if this undefended outline of a view succeeds in avoiding a commitment to the particular "gives rise to" relation of Smith's QWA, the big picture is still radically incomplete. In particular, there are three notable gaps: (1) a full account of what makes some fact the right sort of fact to defeat a *pro tanto* excusing fact, (2) an account of how to get degrees of culpability out of the Sparse Theory, and (3) a clear cut analysis of what goes on in cases of culpable

ignorance. Something will need to be said about each of these if the conjunction of the Sparse Theory and QOW Excuses is actually to deliver on the promise of meeting the liberal's challenge. We will consider each of these issue in turn, but it will be helpful to first consider an objection.

#### 5. THE NATURE OF CULPABILITY

I said that there are two parts to a theory of culpability. The first part was to give an account of what grounds culpability. I offered such an account and supplemented it with a somewhat rough theory of excuse. To answer some of the remaining challenges I mentioned at the end of the last section, I will need to say something about the second part of a theory of culpability. The second part of a theory of culpability is the part that tells us what it is for an agent to be culpable. While I would not aim at something quite so ambitious, we can think of this second part of a theory of culpability as the part that gets at the essence of the concept or the part that provides a real definition of culpability. This task is subtly different than explaining why or when a given individual is culpable.

We should distinguish between two big-picture models of what it is to be culpable. On the *actor-focused model*, facts about culpability are fundamentally facts about the moral status of the wrong-doer or their character; culpability is a kind of moral stain on one's soul or a bad, dark mark in the ledger of one's moral character. According to the *reactor-focused model*, culpability is fundamentally about how other agents should respond to the actions of the wrongdoer in question. That is, to be culpable is to be the appropriate object of another person's blame responses.<sup>13</sup>

There is room for both the actor-focused model and reactor-focused model in our best moral theory, for we can simply countenance more than one concept of culpability. This opens up a possible concession to the liberal. It could be that their view is more plausible on one model and that the moderate's view is more plausible another. The liberal would get something right if this were true, but so would the moderate. My own view is that the moderate's view sits well with the reactor-focused model. Insofar as that is true, there is some important concept of culpability to which the moderate can appeal in stating their view. That is good enough for the purposes of implementing my strategy.

13 These two models go by different names throughout the literature on culpability. The terms I introduce here are my own, which I prefer because they bring out the contrast between the two positions more clearly than other terminology does. For a helpful discussion of how various philosophers have thought about the two models, see the earlier chapters of McKenna, *Conversation and Responsibility*.

And this is so even if we conceded that the liberal's view looks more plausible on the actor-focused model, for it secures the claim that the moderate is on to something deep and important about the nature of moral culpability.

Alright, let us get a toy reactor-focused model on the table. We can start with:

*FA-Culpable*: *S* is culpable for an action if and only if there is reason to have the appropriate reactive attitude toward *S* in virtue of.<sup>14</sup>

A few clarifications: I use the phrase "appropriate reactive attitude" merely to remain agnostic about whether the blame-constituting attitude is resentment, anger, disapproval, some combination of these, or some other attitude or combination of attitudes. Moreover, in the spirit of the Sparse Theory, I will assume that we have a *pro tanto* reason to have the appropriate reactive, blame-constituting attitude toward a given individual only if that person has acted wrongly.<sup>15</sup> The fact that someone has done something wrong plays the role of being a *pro tanto* reason for having the appropriate reactive attitude toward them in many if not most cases. However, this reason can be defeated. Therefore we should think that:

*C-Defeat*: There is a *pro tanto* reason to have the appropriate reactive attitude toward an individual in virtue of an action if and only if is wrong and this reason is not undercut by other morally relevant considerations.

Though it is stated as a bi-conditional, we need not read this claim as a reductive or real definition. The recursive part of the claim therefore should not really bother us too much.

And finally we can plug this all directly back into the Sparse Theory:

*The Sparse Theory\**: The fact <there is a *pro tanto* reason to have the morally appropriate reactive attitude toward *S* in virtue of an act *A*> is explained by the fact that <act *A* is objectively wrong>.

14 This statement takes inspiration from the kind of neo-Strawsonian approach in Wallace, *Responsibility and the Moral Sentiments*. This statement may need to be amended further. For example, perhaps not just any old reason will do. If an evil demon threatens to blow up the world if you do not have a particular reactive attitude toward Felicity when she has done nothing wrong, you have reason to have that reactive attitude. However, intuitively, this does not mean that Felicity is culpable. It is possible, then, that we may need to add a further constraint like the reason in question makes a particular attitude "fitting." I take it that any reactor-focused model will require this sort of caveat, and this amendment and others like it are not an *ad hoc* fix for the moderate view *per se*. After all, reactor-focused liberals and conservatives would face cases like this as well. I thank an editor of this journal for pressing me on this point.

15 The appropriate attitude constituting praise or admiration plausibly has nothing to do with acting wrongly. Thus, acting wrongly is only relevant to culpability. Thank you to an editor for pushing me to clarify this point.

And we can understand QOW Excuses more precisely as follows:

QOW Excuses\*: Some fact *F* is an undercutting defeater for our reasons to have the morally appropriate reactive attitude toward *S* in virtue of a wrongful act *A* only if

1. *F* is about some motive with morally acceptable content that give rise to *A*; and
2. There is no further fact *G* that defeats the exculpatory force of *F*.

With this version of the Sparse Theory and QOW Excuses in tow, we can consider the three remaining gaps in the Sparse Theory\*: (1) an account of what makes some fact the right sort of fact to defeat a *pro tanto* excusing fact, (2) an account of how to get degrees of culpability out of the Sparse Theory, and (3) a clear-cut analysis of what goes on in cases of culpable ignorance. In the next several sections I address these three gaps in order.

Before we move on, however, let us consider why excuses are best understood as undercutting defeaters.<sup>16</sup> The main alternative within the framework I have been developing would be that excuses are rebutting defeaters. Rebutting defeaters can be thought of, roughly, as reasons that outweigh other reasons. If there is a rebutting reason against having the morally appropriate reactive attitude toward an agent in a given case, this would not imply that we lack a *pro tanto* reason for having the morally appropriate attitude toward an agent in virtue of their wrongful action. We would simply lack all-things-considered reason to blame the person in question. However, for many excuses, it seems false that I have any reason whatsoever to blame someone. To illustrate: suppose Erwin presses the doorbell of his friend's house, which he has pressed a number of times before. Little does Erwin know that the doorbell has recently been wired to trigger an explosion thousands of miles away, which will kill dozens. Erwin certainly is not acting from a bad QOW. How could he be? Erwin has no clue that the doorbell has been wired to trigger an explosion. No reasonable person could foresee this. Do we have any reason—even one that is outweighed—to blame Erwin? Plausibly not. In a purely objective sense, however, Erwin acts wrongly. The best thing for the advocate of the Sparse Theory\* to say, then, is that whatever excuse is operative in Erwin's case does more than simply outweigh the reason we have to blame Erwin (in virtue of his objectively wrong action). Indeed, plausibly there is no such reason to blame Erwin. But how could this be?

16 My use of the distinction between “undercutting” and “rebutting” defeaters follows the generalization of the distinction from epistemic reasons to all normative reasons proposed in Schroeder, *Slaves of the Passions*, ch. 7. I thank an associate editor of this journal for pushing me to address my appeal to undercutting defeat.

Undercutting defeaters work differently than rebutting defeaters. Undercutting defeaters “remove” or “disable” our reasons that might otherwise exist were those defeaters not present. Erwin’s excuse, plausibly, is an undercutting defeater. After all, this would explain why Erwin’s act would meet the grounding condition of the Sparse Theory\* and yet there is no reason to blame Erwin. The reason that would otherwise be grounded has been undercut or removed by Erwin’s excuse. In what follows, I will often use the term “defeater” to mean “undercutting defeater” in particular.

## 6. THE CONCERN CONSTRAINT

I aim to defend the moderate against the liberal’s challenge. I have sketched a theory of what grounds culpability and a theory of excuse. This theory of excuse, if it is to be useful for the moderate, needs to explain why some facts count as defeaters for *pro tanto* excuses and why others do not. If this cannot be done successfully, then the moderate is still in trouble. I will not attempt to settle the question of what distinguishes the relevant facts. I will merely attempt to show that the moderate has something plausible to say.

Philip Robichaud and Jan Willem Wieland have offered the beginnings of a response to the liberal’s challenge. They divide the task into two papers.<sup>17</sup> In one, they argue that it does not follow from Smith’s articulation of the liberal’s challenge that blame fails to transfer across the morally relevant relations.<sup>18</sup> If by “follows” they mean “deductively follows,” then surely they are right. But as far as I can tell, that fails to cut to the heart of the liberal’s challenge. The liberal’s challenge is explanatory. It therefore requires an abduction. And such forms of argument are perfectly good even if they are not deductive.

Of course, this still leaves their second essay. In that essay, Robichaud and Wieland defend the following principle:

*Concern Constraint:*  $B_1$  transfers to  $B_2$  only if the benighting act expresses a deficit of concern for the same consideration in virtue of which the unwitting act is wrong.<sup>19</sup>

$B_1$  and  $B_2$  represent the culpability for the benighting act and unwitting act respectively. The Concern Constraint embodies a pretty plausible way to restrict the proper scope of the Transfer Model. And Robichaud and Wieland argue as much.

17 See Robichaud and Wieland, “Blame Transfer” and “A Puzzle concerning Blame Transfer.”

18 See Robichaud and Wieland, “Blame Transfer,” 296.

19 See Robichaud and Wieland, “A Puzzle concerning Blame Transfer,” 17.

But the Concern Constraint does not supply the sort of explanation needed in response to the liberal's challenge. The liberal wants a theory-neutral explanation of why the Transfer Model should be accepted. The worry is that, without such an explanation, the Transfer Model just looks like a dressed up restatement of the moderate intuition. What Robichaud and Wieland have provided us with is the version of the Transfer Model that is the least susceptible to counterexample. While helpful, this is not an explanation of why we should accept the Transfer Model in the first place. Once again, consider Smith:

Of course, it is true that at an earlier time, the time of the benighting act, the agent had a reprehensible configuration of desires—a configuration that typically included a willingness to risk eventual wrong—doing of exactly the sort exemplified in the unwitting act. But the fact that he earlier had faulty motives does not show that he now has faulty motives.<sup>20</sup>

I do not read this passage as pinning the moderate with a counterintuitive implication about who is culpable. It is an attack on the quality of the moderate's explanation of their view. And even if I am wrong about the exegetical point, not much changes. I have simply misinterpreted the liberal challenge in a constructive way: the moderate now faces a new explanatory challenge to which Robichaud and Wieland do not respond.

Robichaud and Wieland might reply that Smith does provide counterexamples. This is true, but we must be careful. Smith employs counterexamples, in my reading, to show the inadequacy of a variety of proposed explanations of the Transfer Model. She does not provide counterexamples to the Transfer Model as such. Her objection to the Transfer Model—and therefore the moderate view—is the aforementioned explanatory objection.<sup>21</sup>

Robichaud and Wieland might reply to this point by claiming that what they did was offer an explanation of the Transfer Model that was not open to counterexample. They therefore did provide a response to Smith. However, they would be wrong to make this response. The Concern Constraint is aptly named. It reads like a constraint on candidate blame transfers, not an explanation of why blame would transfer from an earlier event to a later event. Moreover, when Robichaud and Wieland offer an intuitive gloss of the Concern Constraint, they actually assume what needs explaining:

If blameworthiness is to transfer from  $B_1$  to  $B_2$ , then there must be a match between the kinds of reasons for which the agent shows

20 See Smith, "Culpable Ignorance," 559.

21 Smith confirms this in correspondence.

diminished concern in the benighting act and the kind of reasons that underwrite the wrongness of the unwitting act.<sup>22</sup>

It is the antecedent of their conditional that the liberal is attacking, not the consequent or the conditional itself. The question is why we should be committed to the idea that culpability transfers. Again, we cannot answer this question by simply pointing to our intuitions about culpably ignorant agents because the issue of whether culpably ignorant agents are culpable for their benighting acts is what is at issue.

Here is a different response that Robichaud and Wieland could make. They could follow Daniel Miller's recent suggestion that

an agent's degree of blameworthiness for some action (or omission) depends at least in part upon the quality of will expressed in that action, and an agent's level of awareness when performing a morally wrong action can make a difference to the quality of will that is expressed in it.<sup>23</sup>

Miller's suggestion seems to make progress insofar as it allows us to point to a mental state possessed by culpably ignorant agents at the time they perform their unwitting wrongful actions. Given that Anne is at least nonoccurrently aware that she may have missed some lifesaving information when she performs adult CPR on the infant, then perhaps this counts as a lack of concern.<sup>24</sup>

It would be strained to claim that Anne's nonoccurrent awareness in any meaningful sense "gives rise to" her giving the infant adult CPR. While the state is present in some dispositional sense, it hardly seems to be giving rise to much of anything. So it is unclear whether the mere dispositional presence of this lack of awareness fully explains why the liberal challenge has been met. The general problem is that the liberal has a view about the explanatory role a mental state needs to play. The problem is not about the type of mental state or its content.

Miller might push back. He might say that the non-occurrent awareness is expressed in an "indirect" way.<sup>25</sup> Robichaud and Wieland, similarly, distinguish between "distal" and "direct" motives, writing:

22 Robichaud and Wieland, "A Puzzle concerning Blame Transfer," 15, emphasis added.

23 See Miller, "Circumstantial Ignorance and Mitigated Blameworthiness," 34. Thank you to an anonymous reviewer for bringing this worry to my attention.

24 One salient concern is what we should say about agents who forget about their benighting acts. Miller offers some response in "Circumstantial Ignorance and Mitigated Blameworthiness," 38–39. It is unclear, though, whether his response would apply to Anne were she to forget about sneaking out of the lifeguard training.

25 See Miller, "Circumstantial Ignorance and Mitigated Blameworthiness," 37–38.



We agree that transfer is problematic if, following Smith, *S* is blameworthy for the unwitting *A* only if *A* expresses a deficit of concern on *S*'s part at the time of *A* (which constitutes a direct motive)... [The Concern Constraint] merely requires that the benighting act that led to *A* expresses a deficit of concern (which constitutes a distal motive).<sup>26</sup>

Thus, Robichaud and Wieland—and plausibly Miller—may be unimpressed by my table pounding about the explanatory impotence of Anne's non-occurrent awareness. So long as there is a distal motive expressed in a culpably ignorant agent's unwitting wrongful act, their non-occurrent awareness need not give rise to that unwitting wrongful act.

Which motives count as “distal” versus merely in the past and unrelated? Robichaud and Wieland seem to think that it is (1) those motives that meet their concern constraint and (2) those suitably related to the expression of a deficit of concern in question. But I cannot see how Anne's attempt to save the drowning infant's life “expresses” a deficit of concern. It is not intuitively obvious and I am not willing to take it on faith that such an expression occurs. Perhaps there is some non-question-begging reasons to affirm that such an expression occurs. However, Robichaud and Wieland do not offer any.<sup>27</sup> Any attempt to provide such reasons seems to reintroduce the liberal's challenge: How do we get an “expression” of the intuitively relevant kind of concern without the “gives rise to” relation?

In summary, Robichaud and Wieland's defense of the Concern Constraint is an admirable contribution to those of us who wish to defend the moderate view. I will appeal to the Concern Constraint myself later on. However, it is not by itself an answer to the liberal's challenge. In fairness to Robichaud and Wieland, it is unclear whether they would claim to have offered a full reply to the liberal's challenge. Perhaps they never intended to. In that case, we may well be allies for the purposes of this paper since we both would then recognize that the moderate needs to do more work.

The Concern Constraint may help us do that work. We can adapt it so that it fits in with the conjunction of the Sparse Theory\* and QOW Excuses. Here is a stab at it:

*Concern Constraint\**: A fact *F* defeats a putative excusing fact for a given unwitting wrongful act *A* only if *F* is about a benighting act that

26 Robichaud and Wieland, “A Puzzle concerning Blame Transfer,” 24. Thank you to an anonymous reviewer for pressing this response on behalf of Robichaud and Wieland.

27 Robichaud and Wieland do argue that one reason Smith provides for accepting a timing constraint can be diffused; see “A Puzzle concerning Blame Transfer,” 24. However, this does not constitute an explanation of why they are licensed to claim that the right kind of expression occurs in Anne's action.

expressed a deficit of concern for the same consideration in virtue of which *A* is wrong.

And we might then reformulate QOW Excuses\* as follows:

CC Excuses: Some fact *F* is a defeater for our reasons to have the morally appropriate reactive attitude toward *S* in virtue of a wrongful act *A* only if

1. *F* is about some motive with morally acceptable content that gives rise to *A*; and
2. There is no further fact *G* about a benighting act that (i) expresses a deficit of concern for the same consideration in virtue of which *A* is wrong and (ii) at least partially in virtue of *F* now obtains.

The moderate may wish to eventually generalize 2 so that it sounds a bit less *ad hoc*. By invoking the term “benighting act,” the condition admittedly sounds as if it is crafted to help with cases of culpable ignorance in particular. I, of course, doubt that 2 actually is *ad hoc*. All 2 tells you is how to distinguish between which facts are and which facts are not second-order excuse defeaters—defeaters for putative excusing facts. One would just want an economical way of eliminating “benighting act” language.

One might be worried that clause ii of condition 2 reintroduces the “gives rise to” relation, which I have claimed is the source of the Liberal’s Challenge. This is not so. The “gives rise to” relation was a relation between a motive and a wrongful act. The “in virtue of” relation I discuss here is a relation between two kinds of facts. One kind of fact, *F*, is a putative excuse. The other kind of fact, *G*, is about a motive relating to *F*. All *G* needs to do is partially explain why *F* obtains in a metaphysical sense while citing a deficit of concern. And this is because *G* merely explains why *F* fails to be a good excuse. *G* does not explain why a given agent is culpable for *A*. The wrongness of *A* primarily explains why the given agent is culpable, and—along with the fact that there is no relevant *F*—it thereby fully explains why the given agent is culpable. Nothing about this story requires that the moderate believes that *G* gives rise to *A*. Since such facts do not enter the moderate’s explanatory story, the moderate needs not explain such facts. In other words still, *G* explains why *F* is a bad excuse, not why a given agent is culpable for *A*.

I take it that CC Excuses fills our first gap. It tells us what makes some fact the right sort of fact to defeat a putative excusing fact. Or at least, it is a reasonably plausible first stab. It shows that my strategy to evade the liberal’s challenge by offering an alternative package of views about the nature and grounds of culpability is not dead on arrival.

## 7. DEGREES OF DEFEAT

Let us now turn to our second gap. The moderate needs an account of how to get degrees of culpability out of the Sparse Theory\*. But why does the moderate need an account of how to get degrees of culpability out of the Sparse Theory\*? Well, because they believe that culpably ignorant agents are often partially, but not fully, culpable for their unwitting wrongful acts. And for all I have said so far, the Sparse Theory\* and CC Excuses do not look to accommodate degrees of culpability. We will therefore need to offer some further revisions to the theory. (Notice, of course, that conservatives could appeal to everything I have said in defense of the moderate thus far, but not care about offering an account of degrees of culpability. If no good theory of gradable defeat is found, we therefore are not forced to the liberal view.)

We can supply a gradable theory of culpability by offering an account of partial defeat. There is more than one way to do this. We could, as others have, invoke some form of non-monotonic logic in order to capture the gradable structure of defeat.<sup>28</sup> But since our task is sufficiently simple, we need not deal with the intricacies of non-monotonic logics. Instead, we can build up our theory of partial defeat with a bit of simple math and the idea that degrees of culpability ebb and flow with the strength of the reasons we have to blame others for their wrongful actions.<sup>29</sup>

We can start by determining the strength of a given reason,  $R$ , to have some morally appropriate reactive attitude toward an agent in virtue of their wrongdoing. This reason can be formally represented by a tuple:

$$\text{Reason} = \langle F, A, T \rangle.$$

$F$  is a fact that stands in favor of an agent,  $A$ , performing action of act type  $T$ . (It is worth noting that formally representing reasons with this tuple need not commit us to the claim that reasons are tuples. That is a metaphysical thesis that is likely false and we do not have the space to discuss.) The strength of a reason is a function of the members of said tuple. Let us represent this with:

$$\text{Strength of Reason} = S_{\langle F, A, T \rangle}.$$

We might think that  $S_{\langle F, A, T \rangle}$  is also in some way a function of the degree of wrongness or badness of the action performed by the agent. Next, we can multiply  $S_{\langle F, A, T \rangle}$  by a defeat function,  $D_{\langle F, A, T \rangle}$  for the given reason where

28 See especially Horty, *Reasons as Defaults*; and Bonevac, "Defaulting on Reasons."

29 For an alternative model of defeat that relies only on orderings, see Schroeder, *Slaves of the Passions*, ch. 7.

$1 \geq D_{\langle F, A, T \rangle} \geq 0$ . This allows us to then claim that the degree to which an agent is culpable for an action  $T$  can be represented as follows:

$$\text{Degree of Culpability} = S_{\langle F, A, T \rangle} \times D_{\langle F, A, T \rangle}.$$

When a defeater fully defeats the given reasons,  $D_i = 0$ . If we have defeaters for defeaters—as I have proposed we should in the case of culpably ignorant agents—we can understand  $D_i$  as the following function:

$$\text{Strength of Defeater} = D_{\langle F, A, T \rangle} = 1 - D_{D_{\langle F, A, T \rangle}},$$

where  $1 \geq D_{D_{\langle F, A, T \rangle}} \geq 0$ . Here  $D_{D_{\langle F, A, T \rangle}}$  is the strength of a defeater for the defeater  $D_{\langle F, A, T \rangle}$ . We can determine the strength of  $D_{D_{\langle F, A, T \rangle}}$  by the same function. We can continue to iterate this embedded function *ad infinitum* if need be; the model therefore embodies a kind of recursive structure whereby the strength of each defeater depends on, *inter alia*, the strength of the further defeaters. Consider an example. Suppose  $D_{D_{\langle F, A, T \rangle}} = 0$ . If this is true, then the strength of all of the defeaters for  $D_{\langle F, A, T \rangle}$  will equal 1. That is,  $D_{\langle F, A, T \rangle}$  will equal 1. But if  $D_{D_{\langle F, A, T \rangle}}$  equals 0.3, this is because the defeaters for  $D_{\langle F, A, T \rangle}$  equal 0.7. We can therefore see how this function captures the recursive structure of defeaters.

Let us throw one last widget into our model just to show how malleable it is. We might think that sometimes more than one fact is relevant to an agent's decision on whether to perform some action. Moreover, it could be that the agent is ignorant of several facts. And she could be culpable for her ignorance of each fact to different degrees. How would we represent that? We can represent this simply by modifying the view as follows to account for multiple defeats for a reason,  $\langle F, A, T \rangle$ :

$$\text{Degree of Culpability}^* = S_{\langle F, A, \Phi \rangle} \times \prod_{D_i \in G} D_i.$$

$G$  is the set of all defeaters for  $\langle F, A, T \rangle$  and  $D_i$  represents each member defeater of  $G$ . The function takes the product of the strength of all such defeaters and then multiplies this product, which will be between one and zero, by the strength of the reason under consideration. The recursive structure of each defeater in the set is maintained as long as we claim that the strength of each defeater in  $G$  is determined as we suggested before:

$$D_i = 1 - D_{D_i}$$

There are surely further complications to consider and further ways to tweak the model. For the purposes of this essay, we only need a simple version of the model. The important point is that we can get a gradable structure out

of thinking about culpability in terms of the strength of reasons and partial defeaters that affect these strengths.

Of course, plenty of further, more substantive questions remain. Most importantly: With which gradable properties do the degrees of defeat ebb and flow? This may well be a subject for an entire book. The liberal, conservative, and moderate all need an account of how the strength of excuses might affect how culpable an agent is. That is, some excuses are better than others and therefore exculpate to a greater degree than others. Since everyone needs an account of how to determine the strength of excuses (and therefore first-order defeaters for culpability), I will set this question to the side. But the moderate needs something more: a substantive account of how to determine the strength of those defeaters that undercut excuses. If there were no plausible story on offer, then the moderate would be in trouble.

Luckily, there are plausible candidates. Reconsider Anne. Suppose Anne justifiably had a credence of 0.001 that important, lifesaving information would be shared during the time that she sneaked out for a smoke. Anne would be culpable to some degree (according to the moderate), but it is plausible to think that she would have been more culpable had her credence been 0.5. To wit, the strength of the defeater in this case may be proportional to the degree of undue risk Anne ran by performing her benighting act. Why would the strength of a defeater be proportionate to the undue moral risk run by the agent in performing their benighting act? We might think that our reasons to blame Anne for her benighting act are proportional to the risk she ran. As her actions become riskier, we will have stronger reasons to blame her. And it is then plausible to generalize from our point about risk to say:

*cc Defeater Strength:* The strength of a given defeater for a putative excuse is proportionate to the strength of the reasons we had to blame the agent for their benighting act.

More work would need to be done in order to show that this would work as a general principle, but it does point us in a nice direction. Determining the requisite strength of a defeater by appealing to facts about the strength of reasons we have to blame agents for their benighting acts looks to be nicely principled.

I take it that by now I have done enough to fill the second gap in my alternative QWA of culpability. I have shown how we can think about the structure of partial defeaters in our theory of culpability. This brings us considerably closer to a sufficiently plausible theory of culpability—and one that side-steps the explanatory problems the moderate would incur were Smith's QWA the only account on offer. We can now turn to the final gap to fill in the theory.

## 8. THE APPRAISAL SYMMETRY

The third and final gap that needs filling is a clear-cut analysis of what goes on in cases of culpable ignorance. One might think that this is easy enough to supply given all of the machinery I have laboriously laid out in the preceding pages. The simple version of the story goes like this. Anne performed a wrong action by failing to save the drowning infant. This would usually make her act culpable. However, she had a putatively acceptable motive for her unwitting wrongful act: she wanted to save the child's life. This fact about Anne's motivations is a putative excuse and therefore—absent other defeaters—undercuts our reasons to blame Anne. However, there is another undercutting defeater in play: Anne's unwitting wrongful act is due, in part, to her taking an undue moral risk that was wrong for the same reasons that her unwitting wrongful act was wrong. This second-order defeater partially undercuts the exculpatory force of Anne's QOW excuse. Anne's QOW excuse then only partially defeats our reasons to blame Anne in virtue of her unwitting wrongful act. Anne is therefore culpable to only some degree and the moderate's view has been secured by my model.

Of course, we can analyze the case in this way within the framework I have developed. But what reason is there to think that this is a good way of analyzing the case of culpably ignorant agents? I think that the analogues of the Sparse Theory\* and QOW Excuses\* give us a better explanation of widely shared intuitions about a class of praiseworthy agents. Insofar as we think that our theories of culpability and praiseworthiness are more plausible when structurally symmetrical, this will lend non-negligible support to the view I have developed in this essay.

Consider Saintly Jack. Jack began his life with the desire to do the most good he could. Whenever someone was in trouble and he knew how to help them, Jack would rush to their aid. His only regret was that he could not do more good. So Jack set out to find the best way to help others. After a ton of high-quality research, Jack came to the conclusion that he will reliably help the most people if he instills a disposition in himself to always act from selfish motives. And Jack was right. As a result, Jack very reliably performs the right action at every juncture, though from entirely selfish motives. Further suppose that Jack one day comes upon Jill who is drowning in a pond. Jack thinks to himself, "While I would ruin my new shoes, I could get a substantial reward for saving her!" He then jumps in and pulls Jill out of the pond, saving her life. Is Jack praiseworthy?

According to my intuitions: yes—Jack is praiseworthy to at least some degree. I think this intuition can be felt further if you compare Jack to Tom. Tom just acts from selfish motives. He always has and always will. Jack strikes me as being more praiseworthy for saving Jill than Tom would be for saving

Jerry in similar circumstances. Now this need not be by a huge degree. So long as one gets the intuition at all, I am in business.

This makes perfectly good sense on the symmetrical version of my view. Jack acted rightly. *Ceteris paribus*, this should give us reasons to praise him. However, Jack acts rightly for the wrong reasons. The fact that <Jack performs the right action for the wrong reasons> undercuts our reasons to praise him so long as no second-order defeaters are in play. But, in this case, there is a second-order defeater: the fact that <Jack instilled a disposition to act from selfish motives so that he could help more people>. Moreover, we can imagine that the reasons that justify Jack in instilling a disposition to act from selfish motives are the reasons that justify him in saving Jill's life. This second-order defeater at least partially defeats the first-order defeater given by Jack's selfish motives. And this in turn means that our reasons to praise Jack for saving Jill have not been fully defeated. So a sparse QWA theorist about praiseworthiness seems to capture our intuitions nicely in this case.

Let us now contrast the analogue of my view for praiseworthiness with an analogue of Smith's view and see whether it can capture the relevant intuition:

*Smith's Praise QWA*: The fact <S is praiseworthy for performing act A> is explained by the facts that

1. <Act A is objectively right>,
2. <S had a noble configuration of desires>, and
3. <This configuration gave rise to the performance of A>.

It should be immediately clear that Jack fails the second condition. He does not have a noble configuration of desires at the time he is acting. Or, if one prefers to understand Jack as having some kind of noble second-order desires, Jack would still fail the third condition since this second-order desire did not give rise to his action. By stipulation, a selfish motivation gave rise to his action. Jack cannot therefore be praiseworthy to any degree on such a view. I think this will run afoul of most people's intuitions more starkly than analogous cases of culpable ignorance.

But let us be careful: my point is not that the liberal now faces some new challenge. My point is more simply that the framework developed in this essay can be modified to easily analyze cases of praiseworthy agents like Jack. Insofar as a symmetrical account of praise and blame is attractive, my model seems to provide the tools for a plausible, general analysis for appraising all kinds of agents: culpable agents who are culpable due to the origins of their ignorance and praiseworthy agents who are praiseworthy due to the origins of their motives. And with that, we have filled the third and final gap.

I have thus defended the moderate view against its most prominent challenge. I have done so by offering a novel account of how facts about an agent's

QOW can sometimes function in explanations for why agents are excused. This view is compatible with the liberal, moderate, and conservative accounts of culpable ignorance. This suggests a change in the dialectic: if the liberal wishes to argue against the moderate, they need a new objection or to show that the view developed here suffers from some fatal flaw. At least for the time being, the moderate is off of the explanatory hook they have been hanging on for the better part of the last forty years.

*University of Wisconsin–Madison*  
jpgoodrich@wisc.edu

#### REFERENCES

- Bonevac, Daniel. "Defaulting on Reasons." *Noûs* 52, no. 2 (June 2018): 229–59.
- Capes, Justin. "Blameworthiness without Wrongdoing." *Pacific Philosophical Quarterly* 93, no. 3 (September 2012): 417–37.
- Horty, John. *Reasons as Defaults*. Oxford: Oxford University Press, 2012.
- Husak, Douglas. *Ignorance of Law*. Oxford: Oxford University Press, 2016.
- McKenna, Michael. *Conversation and Responsibility*. Oxford: Oxford University Press, 2012.
- Miller, Daniel. "Circumstantial Ignorance and Mitigated Blameworthiness." *Philosophical Explorations* 22, no. 1 (2019): 33–43.
- Robichaud, Philip, and Jan Willem Wieland. "Blame Transfer." In *Responsibility: The Epistemic Condition*, edited by Philip Robichaud and Jan Willem Wieland, 281–98. Oxford: Oxford University Press, 2017.
- . "A Puzzle concerning Blame Transfer." *Philosophy and Phenomenological Research* 99, no. 1 (July 2019): 3–26.
- Schroeder, Mark. *Slaves of the Passions*. Oxford: Oxford University Press, 2007.
- Shoemaker, David. "Qualities of Will." *Social Philosophy and Policy* 30, nos. 1–2 (January 2013): 95–120.
- Smith, Holly. "Culpable Ignorance." *Philosophical Review* 92, no. 4 (October 1983): 543–71.
- . "Tracing Cases of Culpable Ignorance." In *Perspectives on Ignorance from Moral and Social Philosophy*, edited by Rik Peels, 95–119. New York: Routledge, 2016.
- Wallace, R. J. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press, 1994.