

“Grasping” Morality

Zoë Johnson King
Harvard University

1. A vital theoretical posit

Elinor Mason’s *Ways To Be Blameworthy* offers something from which the literature on agency and responsibility could greatly benefit: a distinction between varieties of blame, and correspondingly between varieties of blameworthiness. Mason’s taxonomy helps to straighten out and systematize a broad range of tempting intuitions about blame and blameworthiness that are otherwise difficult to fit together into a single, coherent theory. For example: when someone who acts wrongly is so deeply in the grip of a false moral view that they have no compunction whatsoever about what they are doing, being utterly convinced that it is the right thing to do, *but* they are this way as the result of hideous brainwashing from early childhood, it can be difficult to know how to react to them. On the one hand, they seem a paradigm of blameworthiness – a moral monster, as it were. On the other, there is a clear sense in which it is not their fault that they have their monstrous character and perform their horrific actions, given the brainwashing. And if none of this is their fault, then it is hard to see how they could be blameworthy for it. Similar points hold of someone who treats others in ways that strike us as moral atrocities and is entirely indifferent to the atrociousness of their behavior, recognizing that it is wrong and yet proceeding with a chilling callousness, *but* who is this way because they themselves suffered similar or worse indignities repeatedly during their formative years. We feel torn in two directions when confronted with such agents.¹ Again, they seem paradigms of blameworthiness – they are engaged in self-conscious, remorseless wrongdoing. But, at the same time, their tragic formative circumstances render this behavior all-too-understandable in a way that seems to mitigate their blameworthiness. Enter Mason: we solve these and related puzzles by distinguishing ways to be blameworthy and then simply saying that many people are blameworthy in one way but not in another way.

Mason’s key distinction is between “ordinary” blame and blameworthiness and their “detached” counterparts.² We do not blame people like the brainwashed wrongdoer and abused wrongdoer in the ordinary way in which we blame our nearest and dearest, according to Mason, but rather in a distinct, “detached” way. Ordinary blame and detached blame differ in the conditions under

¹ For the canonical discussion of this feeling of being torn, see Watson (1996).

² The final chapter before the conclusion also introduces a third type of blameworthiness: “extended” blameworthiness. I will not discuss that here.

which they are appropriate – hence the corresponding distinction between ordinary and detached blameworthiness – and differ further in the responses that we can expect when we engage in them. Specifically, Mason holds that ordinary blame applies only to agents who have acted *subjectively* wrongly – or, as she glosses it, who have acted wrongly “by their own lights” (pp.25, 27). Mason thinks this because she holds that ordinary blame involves “acting as a proxy for [the agent’s] own conscience... simply reminding her of what she already knows to be the case” (p.105). The idea is that ordinary blame involves an expectation that, on reflection, our blamee will come to agree with us that she has acted wrongly, and will then feel remorse, become motivated to make amends, and so forth. By contrast, if someone acts *objectively* but not subjectively wrongly, Mason holds that we cannot expect them on reflection to agree that their behavior is wrong. As a result, she says, we can only subject them to a sort of blame that is “detached” insofar as it is something that “we apply... without expecting a response” (p.131). This is quite different from ordinary blame, which, according to Mason, “is a communicative act” (p.106). And, according to Mason, “[w]hat makes it felicitous as a communication is that it reminds the agent of something that she already knows” (*ibid.*).

I have added the qualifier “on reflection” because Mason allows that it may not be immediately obvious to a target of ordinary blame that she has indeed acted wrongly by her own lights. Mason says that the agent must have “background knowledge that [her] act is morally problematic, even if that knowledge [is] not at the forefront of the mind” at the exact time at which she acts (103). And Mason maintains that “[f]or most of us, when we act wrongly without full awareness, it is just that we didn’t bother thinking: it would not take much at all for us to admit that the act was problematic” (p.90). Mason also allows that someone who has acted wrongly by her own lights may nonetheless need to have the wrongness of this action not merely pointed out to her, but carefully explained. Mason says that “it may take some convincing [the agent], or even educating her”, and that in some cases we are ignorant of the wrongness of our acts because “we turn away from hard questions and undesirable answers”, in which case “we can be brought around to seeing that we have been acting wrongly by our own lights, but it may take some patience” (104).³ Nonetheless, Mason insists, for ordinary blame to make sense it must be the case that the agent’s own moral beliefs *do* in fact impugn her action, so that we can point this out to her – perhaps with some argument or rhetoric, some coaxing and cajoling, or the use of some pedagogical skills – and expect her to eventually come to agree. If we cannot expect such a thing (say, due to the brainwashing), then we can only blame in the detached way.

It is in pairing the distinction between ordinary and detached blame with that between subjective and objective wrongdoing that Mason introduces what I take to be one of the book’s most central concepts: that of “grasping” capital-M Morality. The capital letter M designates something that Mason variously refers to as “the true Morality” (pp.19, 31) or “the correct value system” (pp.32, 37), and also as “the values we share” (p.6) or “the broad view we all more or less agree on” (p.36) or “the broad moral system that we share” (p.77). Thus, Mason assumes not only that there are clear common threads running through the diversity we find in people’s actual moral beliefs – on

³ One might think that this is not a mere *reminder*. That’s what I think. I will discuss an alternative construal at the end of the paper.

which I will say more in the next section – but also that this common core is a correct reflection of the moral facts. With these assumptions in hand, Mason then defends an account of subjective obligation, subjective rightness and wrongness, and the subjective 'ought', according to which the content of subjective obligation is just that "An agent should try to do well by Morality" (p.47). She emphasizes that this means that "acting subjectively rightly [and wrongly] is indexed to a specific value system" (p.34) – namely, capital-M Morality, which is supposed to be not only *our* shared moral system but also *the correct* moral system. The notion of "grasping" Morality is then what ties subjective obligation to ordinary blame, on Mason's picture: she holds that "an agent who does not have a grasp of Morality does not have any subjective obligations, and as a result is never praise- or blameworthy in the ordinary way" (p.75). In short, Mason's view is that ordinary blame applies to all and only those agents who have subjective obligations, which is all and only those who grasp Morality. And, on this view, what most fundamentally unifies the cases of people who can be blamed only detachedly – like a brainwashed wrongdoer (see p.121 for Mason's discussion) and an abused wrongdoer (see pp.145-147) – is *not* that they did not act subjectively wrongly, but rather that they fail to grasp Morality. On this view, someone's failure to grasp Morality is what makes it the case that she does not act subjectively wrongly, which in turn makes it the case that she can only be blamed in a detached way.

Notice that this introduces a complication to Mason's "acted wrongly by their own lights" gloss on the nature of subjective wrongdoing. On Mason's complete picture, someone's acting wrongly by their own lights is necessary but insufficient for them to act subjectively wrongly. For it is also necessary that *their* lights be *our* lights, which are also *the correct* lights. It is necessary, in other words, that they grasp Morality. Similarly, Mason's remarks about the communicative function of ordinary blame must be understood in the context of her views that this form of blame applies only to agents who act subjectively wrongly and that agents who fail to grasp Morality have no subjective obligations. When Mason writes of blamers "reminding" blamees of what they already know, and of acting as "proxy" for the blamee's own conscience, this must be understood as the reminder of someone who grasps Morality *for* someone else who grasps Morality *about* Morality's content. Ordinary blame is an activity that can only be engaged in from one grasper to another, so to speak.

Mason often uses the phrase "our moral community" to refer to the set of people who manage to grasp Morality and can therefore engage in the activity of ordinary blame and be an appropriate target of it (see e.g. pp.32, 34, 38, 40, 41, 72, 82, 106, 108, 115, 121-24, 129, 139, 147-50, 153, 157, 176-77). So, to understand exactly who the members of this set are, we must further examine Mason's central theoretical posit: the idea of grasping Morality.

2. So, what is grasping Morality?

Unfortunately, Mason does not say a massive amount either about what Morality is or about what grasping is. That is because she thinks that there is considerable vagueness here, with the result that a precise analysis cannot be given. We hear that Morality is "a broad value system that covers

a spectrum of philosophical and non-philosophical views about rules, principles, and intrinsic goods", and also that it is "the best version of our current value system, a cleaned-up version of common sense morality, the highest common denominator rather than the lowest" (37). But we do not hear anything about how the cleaning-up process works (reflective equilibrium? Some sort of idealization?), nor how the metaphor of the highest common denominator is to be understood (supervaluation over the cleaned-up versions of people's actual beliefs? Something less formal?). So, it is unclear precisely what the content of Morality is – or, in other words, precisely what one grasps when one grasps Morality. Mason also tells us that "Grasping Morality in my sense does not necessarily involve grasping every facet of Morality. So there is vagueness at the boundaries, it is not always clear when someone is 'in' and when they are 'out' of the moral community" (38). Later in the book, she adds that people in our moral community are "the people we interact with, the people we learn from, the people we have personal relationships with" (82). Since those in the moral community are all and only those who grasp Morality, this is suggestive, indicating that Mason assumes that most of the people around her are not vague borderline cases but rather are among those who determinately succeed in grasping Morality.

That's about all Mason tells us, at a general level, about what grasping Morality involves.

We get more clues by way of examples. Mason tells us that "Kantians and Utilitarians disagreeing about trolley cases do not (usually) think their opponent is a moral monster, they take themselves to be disagreeing within a larger framework on which they basically agree" (p.37). I take it that the idea here is that all parties to trolley disputes – the Utilitarians and the Kantians – are within our moral community and count, despite their disagreements, as grasping Morality. Mason then tells us that Huck Finn is a borderline case, somewhere within the vague area, since his instincts seem pretty decent but he believes that it is morally permissible to enslave people (pp.39, 59-60). Later she adds that it can also be indeterminate whether someone who thinks that gay sex is a sin is thereby placed outside the moral community or is simply someone with a "blind spot" (176). However, Mason takes the case of a dictator's son who has been thoroughly indoctrinated, and who now happily spies on his friends in support of his father's corrupt regime, to be a paradigm case of someone who fails to grasp Morality and cannot be blamed in the ordinary way (pp.122, 154-157).⁴ And she says the same about someone who "has been raised in an entirely sexist environment, which of course has shaped the evidence that is available to him", who now "sincerely believes that women are too silly to be allowed any power or responsibility" and who for this reason forbids his wife from reading books (116-117). Concerning this case, Mason says that "there is no point in blaming him in the ordinary way" and that he "stands outside our moral community" (p.117), but that he is not thereby off the hook and is still properly subject to detached blame in light of the fact that he has acted objectively wrongly.

I confess to finding these puzzle pieces difficult to fit together. As far as the examples go, I assume that deeply-ingrained racism, deeply-ingrained homophobia, and deeply-ingrained sexism are roughly as big of a moral deal (so to speak) as one another. So I struggle to see why it might be

⁴ The example is based on Susan Wolf's (1987) example of "JoJo", but Mason adds some twists of her own.

that the latter one places someone determinately outside of the moral community but the former two just make them a vague borderline case.⁵ Moreover, I think that several Kantians and Utilitarians would strenuously resist the idea that they are disagreeing within a broad framework on which they basically agree. And I think that they would have a point; if one person holds that morality is all about the Principle of Utility while another holds that it is all about the Categorical Imperative, then it is hard to see them as fussing over the details of a broad framework that they share. Their broad frameworks seem completely different.⁶ More generally, philosophers disagree about all sorts of things, including fundamental moral principles and accounts of the final bearers of moral value as well as judgments about the moral status of particular actions. Some deny that the correct value system is "a cleaned up version of common sense morality". So, the idea that we all share a broad framework might be a tad optimistic, depending on who *exactly* is included in the "we" and what *exactly* sharing a framework amounts to. Similarly, the idea that "the people we interact with, the people we learn from, the people we have personal relationships with" are all people who grasp Morality may be unduly optimistic. Speaking for myself, I have interacted with and learned from plenty of hardcore Kantians and Utilitarians, and I have maintained personal relationships with people who harbor deeply-ingrained biases that I don't think I can change their minds about.

One thing that is clear is that Mason allows that people can count as grasping Morality – as being determinately *in* the moral community, rather than being a vague borderline case – even though their moral understanding is less than perfect. This is a possibility that Mason acknowledges at multiple points in the book, although she doesn't discuss many examples of it. What she says is the following:

[I]t is possible that an agent could be non-culpably ignorant of some moral fact, without being outside our moral community. How likely this is depends on how coherent Morality is. If Morality is a set of independent principles, it would be very easy to non-culpably miss some. If it is a set of closely related directives, supported by an underlying rationale, then it is much less likely that an agent could non-culpably miss some moral fact (p.41).

[O]n my account... it is possible that an agent could meet the threshold for moral knowledge and yet not know everything about Morality. So it is possible that an agent could be non-culpably ignorant of a moral fact... Given that, in general, Morality is fairly coherent, it would be surprising to come across someone who had missed an important part. However, assuming that it is possible, such an agent

⁵ It might be that Mason thinks that Huck's good instincts *compensate* for his deeply-ingrained racist beliefs, thereby pushing him into borderline territory. If so, one wants to hear more about how compensation works. I see no similar compensating factor in Mason's example of homophobia.

⁶ In our session at the 2022 Pacific APA, during which I delivered a version of this paper and Mason replied, she said a bit more about what she takes the content of Morality to be: she suggested that it is basically all about doing right by sentient beings. I think that some Kantians would disagree with this – they think that morality is all about doing right by *rational* beings. And plenty of non-Kantian ethicists also doubt that sentience is the ultimate criterion of moral status.

would have the simple excuse of ignorance, and would not be blameworthy (pp.89-90).

It is unclear to me how many people fall into this category, since it is unclear how wrong you can be while still counting as grasping Morality rather than being a vague borderline case. But I don't think that Mason can afford to be *quite* as ecumenical on this point as she presents herself as being. That is because, if philosophers and ordinary folk with massively divergent moral views all count as grasping Morality, then Morality cannot be all *that* coherent. The simpler and more unified it is, the more people determinately fail to grasp it, given the actual diversity in our moral beliefs. So, since Mason is committed to the idea that "we" by-and-large all share this framework, I think that the best option for her is not to leave open the question of how coherent Morality is but rather to emphasize that it is not very coherent; that there is an enormous variety of morally significant things, and that different folks each count as grasping Morality in virtue of their being roughly correct about some of these things – though not necessarily the same ones as each other. But, if that is right, then it would not be at all "surprising" to come across someone who grasps Morality overall has but nonetheless missed an important part. To use Mason's own example,⁷ someone could grasp Morality although she sincerely believes that, while it is usually wrong to hurt people, there is a blanket exception for hurt by offense. This is not at all unlikely or surprising. Or, just to throw more examples out there, perhaps someone who grasps Morality overall could still be mistaken about when she has the standing to interfere in someone's personal affairs, or about the conditions under which she is released from a promise, or about who has a right to be provided with what sorts of information, or about which student excuses are good enough to grant extensions on their papers and which are not, et cetera.

Let's say that agents who determinately succeed in grasping Morality, despite having missed an important part, have *gappy grasps*. There are gaps in their grasp, but they grasp nonetheless.

The possibility of gappy grasps creates a problem for certain parts of Mason's view. I spell out this problem in the next section.

3. The problem

Recall that Mason holds the following:

- a) Agents who grasp Morality act subjectively wrongly (when they act wrongly), whereas agents who fail to grasp Morality can only act objectively wrongly.
- b) Someone is blameworthy in the ordinary way iff they act subjectively wrongly and they have no excuse, whereas someone who only acts objectively wrongly can only be subject to detached blame.

⁷ This example is drawn from the work of Nomy Arpaly (2002), though again Mason adds some twists of her own.

- c) Agents who grasp Morality can be blameworthy in the ordinary way, whereas agents who fail to grasp Morality can only be subject to detached blame.

...with (a) and (b) jointly explaining (c).

Now, consider someone who grasps Morality overall and yet has a significant gap in her grasp – the details don't matter as long as the case has the right structure, so just imagine whatever kind of case you like. Suppose that the precise bit of Morality about which your agent is mistaken pertains to an action that she performs and explains why it is wrong, such that she mistakenly believes her action to be permissible when (to repeat) it is wrong. This agent does *not* act wrongly by her own lights. She grasps Morality, so *usually* when she acts wrongly she acts wrongly by her own lights. But not this time; this time her action falls in the gap, and she believes, falsely, that it is permissible. The same will hold on any other occasion when it is the precise bit of morality that she is mistaken about that explains her action's wrongness. To be clear: how often this can happen depends on how coherent Morality is, as Mason acknowledges, but I have argued that she should say that Morality is not especially coherent, in which case this sort of thing presumably happens fairly frequently. Even if Mason doesn't follow me in emphasizing the disunity of Morality, though, she still must hold that cases with this structure *can* occur, since she holds that people can have gappy grasps. To deny the possibility of such a case would be to deny the possibility of gappy grasps, thereby restricting the moral community to those who are morally omniscient. And that is not a very appealing view.

Mason has two options for what to say about cases in which someone with a gappy grasp acts in a way whose wrongness is explained by precisely the bit of Morality that she fails to grasp. Here are the options:

- Option 1. Say that the agent can be blamed for the action in the ordinary way.
- Option 2. Say that the agent cannot be blamed for the action in the ordinary way.

These are the only options.

Mason's discussion of gappy grasps indicates that she favors option (2); she says that "such an agent would have the simple excuse of ignorance, and would not be blameworthy" (*op. cit.*). But, by stipulation, agents with gappy grasps do determinately succeed in grasping Morality. So this verdict is inconsistent with the claim – central to Mason's picture – that agents who grasp Morality and act wrongly are blameworthy in the ordinary way. Since we are talking about people who succeed in grasping Morality, Mason's claim (c) entails that they can properly be subject to ordinary blame. And claim (c) is entailed by claims (a) and (b), which are also central components of Mason's picture. So she cannot give up claim (c) without major revisions to the whole picture. Although Mason herself seems to endorse option (2), then, it looks as though she would be wise not to do so.

Moreover, the consideration that Mason gives in favor of option (2) looks weak. Mason says that an agent in the sort of case I am interested in would not be blameworthy since she "would have the simple excuse of ignorance". Here the notion of a "simple excuse" is a technical term, which Mason introduces earlier in the book by saying that "an agent has a simple excuse when her agency is not impaired, but it is blocked" (p.22), and citing "unavoidable ignorance" and "sudden paralysis" as examples of the relevant sort of agency-blocking factors. But there is no reason to assume that gappy grasps always involve *unavoidable* ignorance. On the contrary, there are usually plenty of things that someone with a gappy grasp could do to fill in the gap and plenty of things that she could have done to avoid developing it in the first place. And gappy grasps are not remotely analogous to paralysis. So it is hard for me to see how the concept of a simple excuse applies in these cases. In addition, Mason invokes the notion of a simple excuse when discussing a case of a man who routinely litters and does not know that this is wrong; about this case, she says that he "does *not* have a simple excuse" because "he would not disown his action" (p.113, my emphasis). I am not sure exactly what Mason has in mind when she writes of disowning one's action. But I gather that we are supposed to think that the litterer would not disown his action either because he did it on purpose or because he does not think that it is wrong (or perhaps both). And all of this applies to people with gappy grasps who act in a manner the wrongness of which is explained by the part of Morality that they fail to grasp; they do what they do on purpose and do not think that it is wrong. So, if the litterer does not have a simple excuse, then, *contra* Mason's explicit claim about the case, it looks as though the agent with a gappy grasp cannot have a simple excuse either. But if they don't have a simple excuse, that gets rid of Mason's stated reason to take option (2).

Suppose Mason takes option (1) instead. This is a better fit with the overall picture defended in the book, according to which someone acts subjectively wrongly and can be properly subject to ordinary blame just in case she grasps Morality. Option (1) is consistent with claims (a–c), which is helpful for Mason.

Recall, though, that Mason also accepts the following:

- d) Agents who act subjectively wrongly act wrongly by their own lights.
- e) In ordinary blame, we are simply reminding the agent of the wrongness of her action.

Option (1) is consistent with claims (a–c). But it is inconsistent with claims (a–e). To see all this, recall first that, since agents in the kinds of cases we are talking about determinately succeed in grasping Morality overall, claim (a) entails that they act subjectively wrongly and claim (c) entails that they can be blamed in the ordinary way. This all seems fine given claim (b), the tie between subjective wrongdoing and ordinary blame. But claims (a) and (d) jointly entail that these agents act wrongly by their own lights. And that is false. On the contrary, as I emphasized earlier, these agents are by stipulation in precisely the circumstances in which they *don't* act wrongly by their own lights (although they do act objectively wrongly). Likewise, claims (c) and (e) together entail that these agents can be subject to a kind of blame in which we are simply reminding the agent of the wrongness of her action. And that too is false. For, again by stipulation, we are dealing with

precisely the cases in which agents with gappy grasps *cannot* be "reminded" of the wrongness of their action because they do not know and have never known that it is wrong. We cannot remind someone of something that they do not know and have never known. So, since ordinary blame is just a reminder – as Mason holds that it is, per claim (e) – it turns out that these agents *cannot* be subject to ordinary blame after all.⁸ Similarly, since subjective wrongdoing involves acting wrongly by one's own lights – per claim (d) – it turns out that these agents do not act subjectively wrongly after all. But claims (d) and (e) are not optional details of Mason's view; they are elucidations of the concepts of subjective wrongdoing and of ordinary blame as she understands them. So, like claims (a–c), Mason cannot give up on them without major revisions to her whole picture. Option (1) is also a tough pill for Mason to swallow, then.

This looks bad for Mason. To summarize: option (2) is clearly inconsistent with the core tenets of her view, as exemplified by claims (a–c), whereas option (1) is consistent with the letter of these laws but not with their intended substantive content, as indicated in claims (d–e). But options (1) and (2) are the only options. So something, somewhere, has to give.

Before saying what I think should give, let me observe that there is another kind of puzzle case in the vicinity that is effectively the opposite of the kind of case that we have been considering. Imagine that someone has several enormous gaps in her understanding of Morality, sufficient for her to determinately fail to grasp Morality overall – that is, not to be a vague borderline case, but to be determinately on the "out" side of the vague area. But imagine that, on one occasion, this agent performs an action that is wrong and whose wrongness is entirely explained by the bits of Morality that she *does* understand. To repurpose one of Mason's examples as illustration: take the guy who fails to grasp Morality in virtue of his believing that women are too silly to be allowed any power or responsibility and suppose that he tortures some kittens, knowing full well that this is wrong. By stipulation, this agent fails to grasp Morality. So claim (a) entails that he can only act objectively wrongly, not subjectively wrongly, and claim (c) entails that he can only be blamed in the detached way. But this guy understands exactly why it's wrong to torture kittens; the part of Morality that evades his grasp has nothing to do with the wrongness of torturing kittens, which we may suppose that he grasps perfectly. And torturing kittens is what he's doing. So why would he not be subject to ordinary blame?

4. A proposal

The problem is that Mason wants to accept all of the following:

- a) Agents who grasp Morality act subjectively wrongly (when they act wrongly), whereas agents who fail to grasp Morality can only act objectively wrongly.

⁸ Mason might think that the stipulations that someone (a) grasps Morality and (b) acts wrongly jointly suffice to show that, deep down, she does know that what she is doing is wrong. This would mean that the agent does act subjectively wrongly and that our blame is not a mere reminder. But it is incorrect. When someone has a gappy grasp of Morality, her moral beliefs do not encompass the bits of Morality that evade her grasp. That's what it is for there to be a gap.

- b) Someone is blameworthy in the ordinary way iff they act subjectively wrongly and they have no excuse, whereas someone who only acts objectively wrongly can only be subject to detached blame.
- c) Agents who grasp Morality can be blameworthy in the ordinary way, while agents who fail to grasp Morality can only be subject to detached blame.
- d) Agents who act subjectively wrongly act wrongly by their own lights.
- e) In ordinary blame, we are simply reminding the agent of the wrongness of her action.
- f) Gappy grasps are possible.

...But she can't, because these claims are jointly inconsistent. If gappy grasps are possible then it is possible for there to be cases in which someone acts wrongly but the wrongness of her action evades her grasp, lying within the gap. These agents do not act wrongly by their own lights and our blame of them is not a mere reminder, notwithstanding the fact that they do grasp Morality. So Mason's picture entails both that they do and do not act subjectively wrongly, and both that they can and cannot be blameworthy in the ordinary way.⁹

Some part of the picture must be revised, then. What is the best thing for Mason to get rid of?

It would be unwise to deny (f), that gappy grasps are possible. As previously noted, this restricts membership in "our" moral community to the morally omniscient – which, I assume, does not include any real people. So this option is distinctly unattractive.

One possibility is for Mason to double down on (e) and (f) at the expense of (a) and (c), insisting that what matters for ordinary blame is that the agent acted wrongly by her own lights and asserting that even agents who grasp Morality overall nonetheless cannot be blamed in the ordinary way on the particular occasions when they falsely believe that what they are doing is permissible. But this would be suspiciously close to the view that clear-eyed akrasia is necessary for ordinary blameworthiness. Mason considers and rejects that view in chapter 4, holding that someone is blameworthy in the ordinary way if she *could* come to see that she is acting wrongly by her own lights, perhaps with a bit of prodding from those around her, even if she doesn't see it right away.

This is one part of the book where Mason raises the possibility of gappy grasps and sets it aside, saying that such an agent "would have the simple excuse of ignorance" (*op. cit.*) and that she is not interested in such cases. But what she says right afterwards might prove instructive. It is as follows:

I am imagining someone for whom it is true that if they had thought about it a little longer, if they had thought more about their dispositions and alertness to red flags,

⁹ Just in case these entailments are not obvious: given the facts of the case as stipulated, that they do act subjectively wrongly and can be blameworthy in the ordinary way follows from (a) and (c) by modus ponens, but that they do not act subjectively wrongly and cannot be blameworthy in the ordinary way follows from (d) and (e) by modus tollens.

if they had tried harder, in other words, they would have seen that this was a problematic act. (p.90)

An important component of Mason's view – with which I am in full agreement – is that we can be blameworthy for failing to try hard enough to do well by Morality and that our failures to try are sometimes more epistemic than practical, consisting in failures to engage in moral reflection. Mason says that in such cases the agent is blameworthy in the ordinary way and their actions are subjectively wrong, although the agent is not aware of this at the time when they act. So, although Mason is clear at this point in the text that she is *not* talking about agents with gappy grasps, what she says might be extended to them in an attempt to secure the conclusion that agents with gappy grasps can also be blameworthy in the ordinary way. Perhaps at least some of these agents could *fill in their gaps* if they just tried harder to think about it properly. (Recall that their ignorance does not seem *unavoidable*, which is why it is hard to see how they would have a simple excuse.) And, if that is so, then perhaps Mason can say that these agents also act subjectively wrongly and are blameworthy in the ordinary way – despite the fact that they do not act wrongly by their (actual) own lights and that our blame is therefore more than a mere reminder. This would be a way of doubling down on (a) and (c) at the expense of (d) and (e).

I don't think that this is the best option for Mason. Recall that, if Morality is not *that* simple and unified, then it is quite easy for people to count as grasping Morality overall – being determinately in our moral community, rather than in the vague area – despite their moral understanding being gappy. When these agents act in a way whose wrongness is explained by the bit of Morality that evades their grasp, I'm sure that there is usually *some* amount and *some* variety of moral reflection that they could engage in and that would lead them to see the error of their ways. So there usually will be some ways in which they could have tried harder that would have led them to realize that what they are doing is wrong. Nonetheless the subjunctive conditional "*if they had tried harder... [then] they would have seen that this was a problematic act*" will frequently be false of these agents. For it will often turn out that, given the agent as she actually is, it would take rather a lengthy amount of moral reflection for her to fill in *this* particular gap and see the error of *these* particular ways. It might turn out, for instance, that this agent would only come to see the error of these ways if she talked to the right people and read the right blogs and op-eds, which would not be the first ones that she reached for even if she were to engage in further moral reflection about the very moral question at issue. Her material circumstances might also need to change substantially for her to be able to engage with the right sources; she might need to acquire a stable source of income so that she doesn't have to work three jobs and has time for moral reflection, for instance. In short, it might well be that there are *some* possible worlds in which our agent engages in further reflection and changes her mind about the permissibility of her actual action, but all of the *closest* worlds in which she engages in further reflection are ones in which she doesn't change their mind. In that case, Mason's subjunctive conditional is false of these agents. So, even if we do say that ordinary blame applies to agents who could easily come to see the error of their ways upon reflection, we do not thereby secure the result that it applies to agents with gappy grasps.

Although this strategy doesn't work, I think that considering what it would take for people with gappy grasps to fill in the gaps can be highly instructive. That is because I think that the part of Mason's picture that should be revised is claim (b): the part that ties ordinary blame to subjective wrongdoing and detached blame to objective wrongdoing. Considering what it takes for people with gappy grasps to fill in the gaps is helpful because it unearths cases in which, though it would be quite difficult for the agent to fill in the gap on her own, others around her can help her to fill it in – and they can do so in part *by blaming her*.

To warm up to this idea, notice that there is something strange about Mason's claim that ordinary blame is "communicative" (*op. cit.* and pp.99, 102-103, 107, 114-115, 121, 125, 139, 147), given her picture of what it involves. Mason insists that if someone is not acting wrongly by her own lights then any attempts at ordinary, communicative blame would be "pointless" (p.114; cf. p.117, 122). But it is a strange picture of communication that holds that attempts at communication only have a point if the content that I "communicate" to you is content that you already know and that I am simply reminding you of, rather than content that I tell you and you come to learn on the basis of my testimony. Communication is normally a process by means of which information is *conveyed* from one party to another. Its point is to convey the information. So it is not rendered pointless by the fact that the recipient of the information does not already have the information; on the contrary it would be more natural to say that communication is rendered pointless by the fact that the recipient *does* already have the information. If blame were communicative in this usual sense then it would be a process by means of which we *tell* agents who have not acted wrongly by their own lights that they have, in fact, acted wrongly, and through conversation they slowly come to agree with us that they acted wrongly – evidently changing their "lights" in the process. In such cases some information about the moral facts is conveyed by one party who possessed it all along to another party who initially did not possess it, in much the same way that any other information can be conveyed by communication. I do think that blame can work like this. And I think that these cases of blame are among the ones that have the most "point", since I think that blame can serve an important function as a tool for provoking individuals' moral development.

So, return to our agent with a gappy grasp who acts in a manner the wrongness of which evades her grasp. I do not think that it is correct to say that such an agent acts wrongly by her own lights. She grasps Morality, so her lights are generally beamed in the right direction. But the action that she is currently performing remains shrouded in darkness – and perhaps robustly so. Still, I do not think that any of this means that ordinary blame is pointless. And I certainly do not think that the only form of blame we can intelligibly apply to such agents is the detached form that involves no expectation of any kind of response. On the contrary, I think that there is a perfectly ordinary sort of blame conversation that we can have with these agents wherein we *do* expect (or hope, at least) that by the end of the conversation they will come to agree with us. This is because, in *blaming*, we can do something that Mason doesn't discuss at all: we can *try to change the other person's mind*. And this is not futile, so long as there are some kinds of reflection that would suffice to change their minds, even if these are not the kinds of reflection that the agent would most naturally engage in if left to their own devices. We do not have to leave them to her own devices. We can prompt or push them to engage in the particular kinds of reflection that will do the trick.

In doing so, though, we are not just reminding them of moral facts of which they are already aware. Instead, we are *telling* them moral facts – or perhaps we are telling them where to look to find the moral facts, or perhaps preparing them for the idea that there are some moral facts here to be found. The process of trying to convince someone that they have acted wrongly is a complex matter that can include a wide variety of different strategies (and sometimes some skillful maneuvering between them). Still, I maintain that this is a perfectly normal component of our ordinary blame practices. When we see that someone who has in fact acted wrongly has not acted wrongly by her own lights, we do not immediately throw up our hands in despair. We often engage in this skillful kind of blame that is quite ordinary, is genuinely *communicative*, and can be a powerful force for social change.

Word count (including everything, even this): 7,270 words.

REFERENCES

Arpaly, Nomy (2002). *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford University Press.

Mason, Elinor (2019). *Ways to Be Blameworthy: Rightness, Wrongness, and Responsibility*. Oxford University Press.

Watson, Gary (1996). Two Faces of Responsibility. *Philosophical Topics* 24(2): 227-248.

Wolf, Susan (1987). Sanity and the Metaphysics of Responsibility. In Schoeman (ed.), *Responsibility, Character, and the Emotions*. Cambridge University Press, pp.44-62.