Book Review

In the moral responsibility literature, it is often said that blameworthiness
presupposes wrongdoing. But there are numerous conceptions of both blame
and wrongdoing. Elinor Mason's *Ways to be Blameworthy* is, to our know-
ledge, the first book-length attempt to spell out what the relevant kinds of
wrongdoing and blame are that would make a strong connection between
wrongdoing and blameworthiness plausible.

After an introductory chapter that provides an overview of the project,
Chapter 2 begins by defending a central 'responsibility constraint' on the rele-
vant deontic notions: rightness and related notions must be action-guiding
and accessible to the agent in a way that grounds praiseworthiness for right
action and blameworthiness for wrong action. The crucial proposal is then that
this constraint can only be satisfied well enough by a notion of subjective
rightness that tells agents to *try to do well by Morality* and a corresponding
notion of subjective wrongness. 'Morality' signifies *true* morality, and in order
to try to do well by Morality, one needs a sufficient grasp of true morality, a
grasp that has to be in the ballpark, but not necessarily perfect.

In Chapter 3, Mason spells out what *trying* to do well by Morality amounts
to. The notion of trying is reflexive: trying requires knowing, at some level,
that one is trying. It also requires taking steps that one believes are most
likely to achieve one's goal. Correspondingly, failing to try requires that one
has in some way been conscious of the aim one is not trying to achieve. Only
those having a sufficient grasp of Morality can try or fail to try to do well by
Morality.

Chapter 4 further defends the connection between subjective rightness and
what Mason calls 'ordinary' praiseworthiness. Such praiseworthiness, she
argues, is subject to a reflexivity requirement corresponding to that of trying.
This view of praiseworthiness is in conflict with 'searchlight' views that re-
quire full awareness of the morally relevant aspects grounding praise. But,
Mason claims, it is also in conflict with views that merely require that one is
motivated by what is intrinsically good. Huck Finn, though worthy of some
esteem for his concern for Jim, lacks the relevant grasp of Morality to be
worthy of ordinary praise.

Chapter 5 concerns the nature of blame and praise. A key distinction here is between ordinary forms of communicative blame and praise and what Mason calls 'detached' blame and praise. Crucially, ordinary moral blame involves a demand that the wrongdoer recognizes the wrong and feels remorse, and thus presupposes that the blamer is entitled to expect this uptake. We are so entitled only when the blamee has a grasp of Morality. Correspondingly, ordinary moral praise calls on the praisee to take pleasure in the recognition of her efforts to do well by Morality. Detached blame, by contrast, is a reaction to some bad action or agential feature that does not aim for recognition from the agent; perhaps an emotional reaction like anger, repugnance, contempt, or disdain, or the downgrading of expectations on the agent. Likewise, detached praise responds to perceived good action by positive modification of relationships or by positive emotions. Unlike ordinary blame and praise, the detached counterparts are not essentially concerned with the blamee's quality of will.

Chapter 6 turns to excuses. Since ordinary excuses appeal to ignorance or lack of control, and since subjective wrongdoing consists in failure to try to do well by Morality given the information and control available, there are no ordinary excuses for subjective wrongdoing. However, there are mitigating circumstances, showing that the agent did not act as badly as it appears. Under unfavourable circumstances, such as depression, trying might merely consist in thinking that one should do something, even if one fails to act on that thought. The chapter also discusses cases where someone tries their hardest to do right but fails because of some bad motivation, such as the person who does her best to control her racist motivation; these characters are subject to ordinary praise for their attempts, but also detached blame for their failures.

Chapter 7 discusses exemptions, applying the machinery from previous chapters to agents who are in the grip of false moral views, agents who have a grasp of Morality but lack motivation, and agents who fail to see the reason-giving force of Morality.

Chapter 8, finally, argues that we can and sometimes should take responsibility in cases where we do not act subjectively wrongly but do badly, either inadvertently or because of implicit bias or our own bad motives. Taking responsibility will characteristically involve feeling remorse (rather than mere agent regret) and being willing to engage in blame conversations reminiscent of ordinary blame, thus accepting what Mason calls 'extended blameworthiness'. Taking responsibility makes sense because it is required by the sort of attitudes essential to certain relationships or to shared community.

One of the main strengths of the book is that it connects a number of issues that have been discussed at some length in separate debates, letting each issue cast new light on the other. Sometimes we would have wished to hear more about a given issue and how Mason's views relate to other views in the literature, but the discussions are already rich and the general picture as

well as its main elements are clear enough to guide further inquiry. In that spirit, we will look closer at two central aspects of Mason's project, raising a few worries along the way.

1. Subjective rightness and trying to do well by Morality

On Mason's view, one's 'genuine' obligations are *subjective*; to try to do well by Morality. One is praiseworthy if one tries, blameworthy if one fails to try. An important part of her motivation for rejecting more objective notions of rightness and wrongness comes from variations on Frank Jackson's famous case of the doctor who has to choose between three medicines, A, B, and C. The doctor knows that A will partially cure the patient. She also knows that one of B and C will completely cure the patient while the other will kill him, but she has no way of finding out which is which before making her choice. On an objectivist (or, more precisely, what Mason calls a 'hyper-objectivist') notion of rightness, it is right to choose the medicine that will *actually* cure the patient. But, Mason claims, for the doctor to try to do what is objectively right, by picking one of B and C in spite of the risks involved, would be to 'fetishize' objective rightness. The (praiseworthy) thing to do is instead to take into account the values at stake as well as the uncertainty involved, and thus to choose A. One possible conclusion from a case like this is that the notion of rightness relevant to blameworthiness and praiseworthiness is 'prospective': the right action is one that balances the risks in the way that a reasonable person would. However, as Holly Smith has argued, because we might be uncertain about which alternative is prospectively best, trying to do what is prospectively best might also involve unacceptable risk-taking (p. 46). It seems that what the doctor really should do in such a situation, and what would avoid blame, is to take such risks too into account.

What form of rightness might an agent aim to realize without fetishizing rightness in the way illustrated by cases like these? One suggestion is that the agent should do what they take to be best given their beliefs about the situation, including beliefs about risks, or do what is subjectively right, or act in accordance with their conscience. But that seems to get the connection to praiseworthiness wrong: when characters with a terrible moral code follow their conscience, they are not thereby fitting subjects of ordinary moral praise. Because of this, Mason suggests that the relevant kind of subjective rightness be restricted to agents who have a sufficient grasp of Morality. It is when such agents—the *M-competent*, we might say—do what is subjectively right that they are praiseworthy, and when they do what is subjectively wrong that they are blameworthy. The M-incompetent fall under more objectivist notions of rightness and wrongness, but such notions have a different use, identifying moral ideals and deviations therefrom rather than genuine prescriptions and conditions on ordinary praiseworthiness and blameworthiness. (As noted in the chapter overview above, Mason further motivates

the restriction of subjective rightness to the M-competent by appeal to the communicative nature of ordinary praise and blame; we return to this below.)

Suppose that subjectivist notions of rightness and wrongness are needed to avoid the fetishism problem and that a restriction of these notions to the M-competent is needed to avoid the problem of terrible moral codes. Then the question remains how these deontic notions should be understood, more precisely. What Mason suggests is that for an agent to do what is subjectively right is for her to be M-competent and do what she believes is most likely to do well by the morally significant features of the world, or, more succinctly, to try to do well by Morality. An assessment of what is most likely to do well by these features will characteristically involve balancing complex considerations, taking into account both empirical and (restricted) moral uncertainty (p. 64). Even without any idea about what action is hyper-objectively or prospectively best, an M-competent agent can try to do well by the values and risks involved, and doing so would not be fetishistic. (This is not to say that the M-competent agent can always be guided by the prescription to try to do well by Morality so as to avoid blame and be worthy of praise. We can clearly be self-deceived about what we are trying to do. But no reasonable account of subjective obligation can avoid the problem of self-deception, and Mason argues that her account gets the connection to responsibility right: we deserve no credit when we mistakenly think that we are trying to do well by morality.)

Someone coming from the literature on moral uncertainty and the debate about objectivist, prospectivist, and subjectivist rightness might expect Mason to offer a general substantive account of what agents really should do, explaining why they should deal with uncertainty in specific ways in cases like Jackson's or Smith's. At times, it might also seem that Mason is offering pieces of such an account. Given that trying to do well by Morality is sufficient for subjective rightness among the M-competent, some things she says about trying might suggest that what it is right to do is mainly up to the agent. For example, she says that an agent 'is trying to achieve a sub-goal if she does what *she takes* to be appropriate to achieve her sub-goal, given the other sub-goals that fall under her overall goal' (p. 65, emphasis added) and that the agent 'counts as trying when she does what is appropriate *by her own assessment*' (fn. 10, p. 69, emphasis added). But Mason also thinks that such assessments can be mistaken, and that there are better or worse ways of balancing various considerations. For example, she says that:

> [s]ometimes doing well by Morality may look more consequentialist, and would involve balancing risk and value, and in that case, that is what the agent *should* be aiming for. Sometimes compromises may not be possible, and doing well by Morality requires doing only the very best acts. In that case that is what she *should* be aiming at. (p. 47, emphases added)

Mason also claims that we should 'seek out more information when *appropriate*' (p. 47, emphasis added) and that '[t]he more important the goal, the more precautions we *should* take, the more contingency we *should* build in' (p. 67, emphases added). If all these passages employed the same normative notions, they would seem to provide conflicting normative guidance. However, a better interpretation distinguishes the requirement of subjective rightness—the requirement *that one tries* to balance the various considerations well—from requirements on success in what one is thereby trying to do. The block quote and the citations following it are concerned with the latter. The idea is then that the agent can make his or her own assessments about what is appropriate and still qualify as trying to act well by Morality as long as she does so within certain limits set by the requirements of success. Mason's analogy between trying to do well by Morality and trying to be a good parent is suggestive – trying to be a good parent means acting in a way that is not too far from *really* being a good parent (p. 70). But it is only a suggestive analogy. We still need to know more about the limits within which we can make our own assessments and still qualify as trying to act well by Morality. Without such knowledge, it will be very hard to decide whether an agent is praiseworthy in the ordinary sense. Of course, Mason could reply that this is as far as we can get and that we just need to lower our expectations. But we are not convinced that we have to accept this more pessimistic conclusion yet.

Another issue about trying to do well by Morality concerns the relevant notion of trying. On page 65, Mason characterizes the relevant *strong* sense of trying in terms of the following sufficient condition: 'An agent is trying to achieve her overall goal if she does what she takes to be most likely to achieve her goal to a sufficient degree'. On its most obvious reading, however, this characterization is in tension with a claimed virtue of the identification of rightness with a species of trying: it is supposed to explain how 'what an agent subjectively ought to do [isn't] determined by beliefs that she could improve right now if she bothered', beliefs that might be 'lazy, hasty, poorly motivated, or self-deceptive' (p. 45). The explanation is supposed to be that *trying* to do something involves being open to information relevant to success in what one is trying to do (pp. 47–8), thus involving a *general* 'attitude of alertness, a readiness to examine our beliefs, a disposition to double check our evidence in the light of our aims' (p. 45). But, intuitively, even if one satisfies the condition given on page 65, one's belief that one is doing what is most likely to succeed might be self-deceived, due to half-hearted trying. In keeping with Mason's overall project, then, it would perhaps be more natural just to say that trying to achieve an overall end is a matter of being governed by one's concern to achieve that end *in a way that trumps other ends*. Being so governed will plausibly ensure that one is doing what one at least implicitly believes is most likely to achieve the end, making that belief a necessary rather than sufficient condition of strong trying. (Suppose that strong trying

is understood in this way. Then Mason's explanation of why lazy ignorance fails to ground subjective rightness and praiseworthiness mirrors a corresponding explanation available to what Mason calls 'moral concern' theorists in the Strawsonian tradition: if an agent cares as much about relevant values as morality demands of her, she will be open to and act on available information relevant to promoting these values.)

We have just seen that for Mason, the prescription satisfying the 'responsibility constraint' is that of trying to do well by Morality, and doing so involves not only having a certain goal but also having a set of dispositions to examine one's beliefs and evidence. But it is not obvious that such a goal and such epistemic dispositions are always available even to an M-competent agent in a situation where there is something that she ought to do. But if they are not, perhaps in situations that call for swift action, Mason's proposal might get in trouble with the idea that what is genuinely prescribed must be accessible, or that 'ought' implies 'can'.

We are not sure how serious this problem is, and it might depend on what trying to do well by Morality is more exactly, and how strong the dispositions are supposed to be. Perhaps the availability of both goal and dispositions can be understood in a minimal way. Perhaps trying to do well by Morality can be understood merely as trying to satisfy some perceived requirements posed by moral considerations. So understood, trying to do well by Morality seems no harder than trying to do anything else in response to reason. And perhaps the dispositions to examine evidence can be understood as special cases of the dispositions that come with trying to do anything, namely dispositions to be concerned with what affects the prospects of success. If so, they add no further difficulty.

2. Grasping Morality and ordinary blame and praise

The restriction of subjective rightness to the M-competent is central to *Ways to be Blameworthy*. Without the restriction, subjective rightness is unmoored from praiseworthiness, and the idea that agents should do what they take to be right seems deeply problematic. With the restriction, the idea seems promising: put generally and in vague terms, we want people who have a good enough grasp of moral considerations to try to do well by these considerations, and we think that it is good that they try even if they might nevertheless occasionally go wrong.

Corresponding to this restriction is a version of the familiar idea that ordinary blame and praise are communicative. On Mason's version of this idea, ordinary praise and blame aim at getting the target to recognize (the recognition of) their efforts to do well by morality, or their failure to put in that effort, and to respond with pleasure or remorse. Because we are entitled to expect communicative success for this sort of praise and blame only if the

target is in a position to recognize whether or not they put in the effort, only the M-competent are fitting targets, Mason argues.

The strategy of restricting ordinary blame- and praiseworthiness with reference to the fitting targets of ordinary communicative blame and praise is both natural and promising. But the general strategy is not new, and we are not yet convinced of Mason's particular way of developing it.

One initial question concerns what counts as communicative success of ordinary blame and praise on Mason's own view. In Chapter 5, the expectation or demand essential to ordinary blame is first said to be two-fold: it is an expectation or demand that the blamee will recognize and accept the blame as well as an expectation or demand that they have appropriate emotional reactions of remorse to this recognition (pp. 103–5). If these are the expectations, it would suggest that ordinary blame is felicitous only when the blamer believes that the target can understand that they have done wrong and respond emotionally to it. Later, however, Mason says that blame and praise essentially involve an expectation that the target will recognize the blame and praise but only a *hope* for an emotional reaction of remorse and pleasure, thus drawing a distinction between the two responses (p. 111). In line with this, Chapter 7 portrays Scrooge as a fitting target for blame merely in virtue of having the capacity to recognize the wrongdoing, even if he is constituted so as to never be moved by Morality (p. 168). Perhaps the latter formulation is Mason's considered view. But it is not clear that the distinction can be derived from the aims of ordinary communicative blame, which characteristically demands more from its target than the unmoved recognition of wrongdoing one can expect from Scrooge.

We are also not convinced that ordinary communicative blame and praise seek recognition of failure and success in *trying* to do well by Morality, as distinct from recognition of significant failure and success in *actually* doing well by Morality, where doing well by Morality is understood as acting from an adequate concern for relevant moral considerations, or as adequate reasons-responsiveness. After all, blame is often directed at significant failures to give adequate weight to specific values, harms or risks, and only sometimes at general moral laziness. Likewise, praise is often intuitively directed at remarkably good responses to specific needs or risks, and only sometimes at the general *effort* to respond in the right way to moral considerations. It is of course also true that when someone who is generally a good person really tries to balance considerations in the right way, praise might be fitting even if they get things wrong. But that observation seems to be straightforwardly accounted for by the moral concern proposal: the agent might deserve ordinary praise for having adequately responded to the complexity and uncertainty of a situation by trying to get things right even if they fail to display an ideal concern for

other considerations, where the latter failure may or may not be significant enough to be blameworthy.

On Mason's overall view, the sort of praise and blame that we have in mind here would qualify as 'detached' rather than 'ordinary' praise and blame. But insofar as it aims at pleased or remorseful recognition of remarkably good or bad reasons-responsiveness by the praisee or blamee, it is not clear why it should be understood as *detached*, and given that it seems as common as praise and blame for trying or failing to try to do well by morality, it is not clear why it shouldn't be understood as *ordinary* blame and praise.

Suppose that ordinary communicative blame and praise do indeed aim at engaged recognition of significantly inadequate or adequate responsiveness to moral reasons. Then that might leave untouched Mason's motivation for restricting ordinary blame- and praiseworthiness to those with a good enough grasp of Morality: only from agents with such a grasp are we entitled to expect engaged recognition of inadequate or adequate reasons-responsiveness. But the restriction would plausibly be weaker than it is on Mason's *trying*-account, according to which the agent must have known at the time, at some level, whether she was trying or failing to try. For even if one did not already understand one's moral responsiveness as actually significantly inadequate or exemplary, one can come to recognize it as such in response to blame or praise. Blame and praise can themselves be revelatory. What matters for the felicity of acts of ordinary blame and praise would then be that the target can recognize the inadequacy or adequacy of her reasons-responsiveness in light of these acts. Correspondingly, it is tempting to say that an agent is the fitting target of ordinary blame and praise if she has been *recognizably* inadequately or adequately reasons-responsive. (In Chapter 8, Mason suggests that agents can 'take responsibility' in some cases like these, and thus become fitting targets of communicative blame that demands remorse from them. But it is not clear to us that such blame would be unfitting in the absence of prior responsibility-taking.)

We are not saying that this alternative way of understanding the aims of ordinary blame and praise is the right one, or that blame- and praiseworthiness should be understood in terms of these aims. However, it seems to us that the book's current argument fails to favour her alternative over this, and we would like to hear in more detail from Mason why hers is to be preferred.

Our aim in this review has been to give readers a general sense of the contents of the book, in part by raising worries for some of its key contentions as we understand them. While we remain unconvinced of the idea that ordinary blame and praise and corresponding notions of rightness and wrongness relate to the notion of trying to do well by Morality, we think that it is a suggestive idea, well worth thinking about for those interested in blame and praise and their connection to deontic notions. The same is true about many less central claims made in the book, and though the

reinterpretations of familiar phenomena invited by the book's framework don't always wear their advantages on their sleeves, they often force one to see the phenomena with fresh eyes.*

Department of Philosophy, Stockholm University   GUNNAR BJÖRNSSON
Lund Gothenburg Responsibility Project
gunnar.bjornsson@philosophy.su.se

Department of Philosophy, Stockholm University   KRISTER BYKVIST
Institute for Futures Studies, Stockholm
krister.bykvist@philosophy.su.se