

Comment on Colin Marshall's *Compassionate Moral Realism*

Patricia Marino

It is my pleasure to comment on Colin Marshall's book, which I found to be an engaging and refreshingly straightforward read. The book has several virtues that not always present in philosophy books, and that strike me as particularly uncommon in metaethics books: these include a willingness to engage what might seem more like "common sense" questions rather than just questions that tend to be the ones discussed by other philosophers, an explicit discussion of the audience for the book and the reasons the problems are thought to be important, and a sophisticated way of mixing the narrower questions the book is addressing in detail with questions that would be beyond its scope. Every book should have a "Fights I Am Not Picking" section and, even more importantly, a "So What" section.

While I do consider myself a metaethicist, my work does not directly engage the realism/anti-realism distinction (a fact that has, on occasion, prompted shock and confusion from some other philosophers). And my comments here will not focus much on the specific aspects of Professor Marshall's discussion that have to do specifically with that distinction. Instead, I want to use this comment to think about a few broader but related topics related to ideas in the book and particularly to its framing, audience, and questions engaged. My comments are not intended to be criticisms, but rather starting points for thinking through some of these related -- and I think complex and important -- questions. The question concern not so much the particular arguments in the book, but more how the ideas relate to other ideas. Very roughly speaking, I'll be addressing primarily the choice of how to frame the central problem or question of the book and the intended audience for the book -- by which I mean not the literal audience but rather the people who may be expected to learn from or be persuaded by its arguments. I then include a

friendly point of agreement about coherence and a suggestion for how incorporating that point might strengthen and support Professor Marshall's overall arguments.

As a way in to some of these questions, let's look first at what kind of realism compassionate realism is. As we've heard in Professor Marshall's discussion, this realism centers on the idea that we perceive others' feelings, and that to be indifferent to, or to enjoy, others' pain is a mistake. I really like the very clear summary Professor Marshall gives toward the start of the book when he says "One way to gloss the central claim of this book is to say that compassion is the perception of others' feelings and that sadistic pleasure and indifference to others' feelings are perceptual lacks. This, I believe, has implications for morality: it answers the question of why we should be moral and, together with some further, fairly modest assumptions, provides the keystone for a grounding of morality that secures it against the threats of egoism, relativism, and other rationalizations people use for cruelty and apathy."

The point of developing this form of realism is to thus address questions like "why be moral" and to ground the possibility of objective -- or at least, non-relativistic -- moral facts or obligations rather than to draw out any particular conclusions about how to be moral or what to do. Professor Marshall emphasizes that his account is compatible with a range of view in normative ethics, and does not take a stand on the debate over consequentialism versus deontology (see Ch 7). Included in the list of issues that are not being engaged in this book are ones over what Professor Marshall calls "controversial moral cases." Speaking of different paradigms that can ground different kinds of obligations, Professor Marshall says "my argument stays neutral on a wide range of controversial moral cases, especially cases that involve a conflict between these different paradigms."

Generally, issues of moral realism and its relation to objectivity and moral action can, I think, be motivated and understood in several ways. One way these issues are sometimes framed in recent philosophical literature is through a more general epistemology and metaphysics lens. For those of us who are inclined toward philosophical naturalism, the question of how there could be normative facts, and how those facts could possibly motivate us, presents a challenge. Over the last fifty to a hundred years or so, we have seen a vast literature devoted to exploring what options we have for navigating this challenge. Maybe moral statements are fictions, and thus factual but false. Maybe moral statements are more like expressions of attitudes, and thus not even aiming at the same status as natural facts. Maybe moral properties are like secondary properties. Maybe moral factuality and truth can be squared somehow with what Mackie thought of as the "queerness" of moral properties.

It's always seemed to me characteristic of this "M and E" metaethics literature 1) that it is primarily focused on solving an internal philosophical problem, rather than looking outward toward implications, and 2) that it is considered normal and even expected that a good answer to the problem will not have any particular implications for normative or applied ethics. With respect to 1): that is, the problem arises largely because we have a set of commitments about the nature of reality and language that do not seem to square, and the aim is to find a way to square them. With respect to 2): this manifests itself partly in how the answers get evaluated for plausibility in further philosophical writing, but also in informal conversations at conferences and so on. I may be particularly attuned to these informal conversations because that is not at all how I do the metaethics that I do. My work on value pluralism, moral dilemmas, and coherence uses a metaethical framing to draw conclusions about normative ethics and applied ethics. I

remember being at a metaethics conference relatively early in my philosophical career and hearing that good metaethics should never have normative and applied implications, and should, in fact, leave everything alone.

Encountering this idea that good metaethics leaves everything along, at first I thought "Uh oh. Am I doing it wrong"? But over time I came to think that this informal view is mistaken -- or, at least, it only applies to the branch of metaethics that I've just been talking about, that is, the "M and E" metaethics, motivated and framed as a solution to an internal philosophical problem. There is still a way to do metaethics that starts from a different set of questions, questions about how to live and how to make moral decisions. Ideally those answers would also fit with what seems epistemically and metaphysically plausible, but the focus would be on changing or newly understanding our moral lives and practice. Ruth Marcus's work on moral dilemmas, which I build on in my own work, exemplifies this second, more "applied" framing. In a series of classic papers, Marcus used modal logic to draw conclusions about dilemmas, then used conclusions about the formal status of dilemmas to draw practical conclusions -- such as her conclusions that moral dilemmas are real, that they reflect a residue or remainder related to the fact that not all of our obligations were met, and that we should therefore try to arrange our lives and institutions so that dilemmas are less likely to arise. That principle is now taken up and used in contexts like bioethics -- where is sometimes called a "mutuality principle," and in applied contexts like engineering ethics.

To me, it is hugely to the credit of Professor Marshall's book that it engages with what I take to be the second framing, the one in which we are trying to figure out what to do. This is not to speak in any way against the more "M and E" kind of metaethics -- it's just that there has been

so much great philosophy done in that tradition, and, at least recently, relatively less in the more practical style. And it is not to say that implications are Professor Marshall's central motivations. But Professor Marshall goes out of his way, it seems to me, to draw attention to the intended audience for his book and to what the ideas in the book should mean for them in a practical everyday kind of way. For example, we've seen that the main point is to "[secure morality] against the threats of egoism, relativism, and other rationalizations people use for cruelty and apathy." Later, Professor Marshall explores the ways that his view can be used to make people see, or understand, or feel, that if they fail to care about morality they are doing something wrong. Acknowledging that radical skeptics can always fail to be moved by arguments, Professor Marshall points out that there is another audience we should try to reach: people who "actually value being in touch with the affective states of the various creatures their actions can affect." Some of these people might be somewhat compassionate, and some might lack compassion, but Professor Marshall argues that his given answer to the "why be moral?" question can "have force for" all of these people. "Reflecting on the answer," he writes, "good people can take comfort in knowing that they are not simply dupes, and might take this as giving them reason to cultivate compassion in others, such as their children." Further, "People who are not fully compassionate might be moved to cultivate their compassion, or at least stop thinking of compassionate, morally good people as dupes." This audience, he says, is important: "Radical punk rock skeptics should not be the only object of philosophers' attention," he writes. "It is at least as important to help strengthen the moral resolve of good people, and to sway the rare person who is on the fence about whether to develop their moral tendencies"

As I have said, I applaud this framing, and I admire the seriousness with which Professor Marshall has undertaken the difficult -- often surprisingly difficult -- task of being explicit about the target audience for his philosophical arguments and how those arguments apply to them, and especially about identifying a target audience beyond the audience of other philosophers working on a particular topic. I am, though, somewhat skeptical of the robustness of the connection between realism, as understood here, and the effect described. Part of my skepticism has to do with my uncertainty about how much people care about being wrong about something. Does the idea they're being wrong about something disturb or threaten people enough to make them rethink their actions?

Philosophers care a lot about avoiding being wrong about something. But as Jonathan Haidt has put it, philosophers are also "trained and socialized" to care about epistemic factors in a certain way. The mechanism through which Professor Marshall articulates the connection is developed through a proposed emotional resonance we may have to avoiding mistakes and knowing the truth. As in the movie *A Few Good Men*, we bristle at the idea that we "cannot handle the truth," and we take pride in being epistemologically on the ball; in the context of the given account of moral realism, this allows for a source of potential pride for morally good agents. I haven't seen the movie. But if someone I disagreed with said charged me with being immorally indifferent to suffering, and focused their critique on my epistemic shortcomings in this way, it's hard for me to imagine being moved by that, if I haven't been moved by the suffering in the first place. Of course, these are just my impressions about a very broad and open-ended set of questions about human nature and moral psychology. But I do think it's a question why a person who doesn't care about others would care about making a mistake.

A further and related question I would pose about the framing and audience has to do with the choice of emphasizing the "why be moral" aspect of this moral realism over the objectivity and potential normative ethics aspect. As Professor Marshall says in the passage quoted above, one of the targets meant to be kept at bay here is ultimately relativism. On the face of it, the claim that pain is objectively bad is compatible with a wide range of moral beliefs and incompatible moral systems -- indeed, as presented here, it seems compatible with the idea that preventing suffering is a crucial and largely stand-alone kind of moral obligation, grounding our ethical sense and appropriate normative ethics, and also with the idea that preventing suffering is just one of a range of moral sources, alongside others like justice, promise-keeping or honesty, as in theories like that of W. D. Ross's principle pluralism. I pause to note here that "pluralism" as applied to Ross means that there is more than one principle; this is different from the "pluralism" about moral knowledge that Prof Marshall discusses in the book which has to do with pluralism about the nature of moral judgments and knowledge.

Going further, the claim that pain is objectively bad seems compatible with the idea that people may care about different values, in different ways, leading to different conclusions. In my own work, I've argued for a view of coherence in which there can be multiple sets of beliefs that disagree with one another yet are still fully internally morally coherent. One way this can happen is when we share values, but prioritize those values in different ways. Suppose a person must lie to keep a secret. Two moral evaluators might both value honesty and fidelity. But the one who values honesty most may conclude that telling the truth is the right thing to do overall, while the one who values fidelity most may conclude that lying and keeping the promise is the right thing to do overall. A person might prioritize honesty and build a fully coherent system based on that,

or they may prioritize fidelity and build a fully coherent system based on that. These two moral frameworks may come to different conclusions about whether to lie, yet each will be fully internally coherent. On the face of it, compassionate moral realism seems compatible with the existence of these multiple differing frameworks, as long as they all include some form of non-malevolence or benevolence -- caring about and acting to reduce pain .

Similarly, In bioethics, a commonly used framework is principlism, in which there are multiple values such as benevolence, non-malevolence, justice and respect for autonomy; when these conflict, as they often do, we have to use our judgment to "balance" them. Since balancing can be subjective, principlism is the kind of methodology that can lead us to a different kind of "pluralism" -- the idea that more than one theory can be correct -- and from there to charges of relativism. It is, after all, pretty subjective and the outcome can vary from one person to another. Again, the moral realism offered here seems compatible with this strategy too, as long as non-malevolence is a factor we all must consider. As Professor Marshall says, "Since CMR only requires that there be some knowable moral facts, it can therefore leave open the question of moral knowledge concerning other (more debated) topics."

Professor Marshall doesn't explore much the question of compatibility with different normative systems, suggesting that the view is orthogonal to questions in normative ethics. Professor Marshall notes that the arguments of the book could be read as if they favored consequentialism, and he goes out of his way to say that this isn't really the case. In a passage I found interesting and illuminating, to show how the grounding of moral realism in terms of perception can be compatible with deontological as well as consequentialist views, Professor Marshall proposes a metaphor in which different forms of perception can function

metaphorically to ground consequentialist and deontological approaches. He says that the conflict here "is like the conflict of whether the earth is more accurately perceived from the ground or from space."

One reason I like this metaphor is that my own intuition is that to understand fully requires both kinds of perception; likewise, I would say that our obligations are sometimes impartial obligations toward everyone and other times narrow and specific obligations to particular people in virtue of the particular relationship you have with them. However, from my point of view, where the morality as perception metaphor becomes challenging is with the phrase "more accurately." The difficult question in morality seems to me to find a way to balance two different kinds of cares. To be a "realist" about this would seem to require a lot more than realism about pain's objective badness; it would require realism about a wide range of values and how they can be compared. I was curious when reading whether Professor Marshall ultimately thinks that his moral realism really is a better fit with consequentialism, or whether he thinks that the grounding for other kinds of obligations are going to have to be very different from those for sparing others from pain. It is not a criticism of the book that it does not engage these questions -- a book can only do so much, and moral realism is a huge topic! But it does seem to me that the more the framing of the question has to do with the applications and implications, rather than just the "M and E" framing, the more these questions are thrust into the spotlight, and the more a reader becomes interested in them.

As someone immersed in the study of coherence, I'd like to make one final and friendly comment on the use of coherence in the appendix discussion. Coherence is important to Professor Marshall's concept of "functioning"; this concept helps explain the connection between

feeling the pain of another and acting to alleviate that pain. The three basic levels of coherence here are "felt motivation," "action," and "comparative intensity." As I understand it, Professor Marshall's main concept of "agential well-functioning" requires only these first three levels. He then goes on, however, to discuss a fourth and "highest" level: full coherence. He writes, "The highest level of well-functioning along these lines would ... involve having no conflicts between [affects, motivations, and actions]. A perfectly coherent agent, for example, would never be simultaneously pleased and pained by the same thing, would never have competing motivations of any sort, would never desire anything she knew was impossible, and would never act in a way that conflicted with her mental states in any way. Arguably, this perfect agential integration reflects one ideal of rationality."

I disagree with some of these claims about what "full coherence" requires -- in fact, I've written several papers defending the idea that being simultaneously pleased and pained by the same thing, and other forms of ambivalence and conflictedness, are not only highly characteristic of human existence, they're also fully coherent and rational ("On Essentially Conflicting Desires; Ambivalence, Valuational Inconsistency, and the Divided Self," "Moral Rationalism and the Normative Status of Desiderative Coherence). This material about "full coherence" is toward the very end of the book, so the discussion is brief, but I get the sense that Professor Marshall may agree with me. Later, he discusses situations in which morally good agents are pulled in more than one direction at once. Examples include situations in which two creatures are in desperate need but only one can be saved, or in which a creature has two equally deep but incompatible desires. Professor Marshall writes, "My own intuition is that feeling torn is the morally ideal response to such situations, so that, in a world like ours that includes such conflicts, a perfectly

virtuous agent could not be perfectly well-functioning. Surprisingly, then, there may be situations in which perfect coherence precludes perfect moral goodness. My overall argument says neutral on this issue, however."

As they say, one person's ponens is another one's tollens. I completely agree with Professor Marshall that feeling torn is the morally ideal response to such situations. But in that case, I am inclined to conclude not that perfect coherence precludes perfect goodness, but rather that perfect coherence does not include the rejection of ambivalence, torn moral attitudes, and so on. This alternative understanding of coherence may be a better fit with the ideas in the rest of the book than what Professor Marshall called "full coherence," as it allows that our response to various people in pain that cannot all respond to should naturally include ambivalence and regret.

If we go in that direction, the line of thought brings out what I think could be an important further implication of Professor Marshall's ideas: a way in which compassionate moral realism is compatible with conflicting obligations and thus possibly with moral conflicts and dilemmas. For various reasons, realism is often associated with the idea of a definite moral answer to every moral conundrum, and this is sometimes seen intuitively as antithetical to seeing conflict and dilemmas as having an important role in moral theory. Professor Marshall's examples remind us that even the straightforward obligations like "act to reduce pain" can engender moral conflicts and torn feelings all on their own.

Of course, if you go far enough in this direction, you can find yourself back with the Ross and the pluralists, with their intuitive balancing about what is overriding, and all the feeling of subjectivity that this brings in its wake. That might even feel like relativism. As I see it, whether to count a view in which pain is bad but there is no single right answer to a range of moral

questions "realism" depends on what you want your word realism to do for you. As I've suggested, if it's to counter apathy or bad behavior, I'm not sure how far that can go, because I'm not sure whether apathy and unwillingness to face facts are things people appeal to when they try to justify to themselves why they think it's OK to act as they do. In my own work, I've explored the possibility that people are often more morally moved by literature, art, stories, and people they know than they are by arguments. Maybe the good people who need a nudge to expand and act on their compassion would be as well served through community and relationships instead of knowing that they are right about something. Again, this is not to criticize the ideas in the book, but more to reflect on the question of what metaethicists are talking about when they talk about "realism."