

It is a pleasure to comment on Brian Weatherson's insightful, comprehensive, rigorously argued book. *Normative Externalism* seems to me an example of one kind of analytic philosophy at its best -- it is a showcase of how to pay careful attention to details, to tease out distinctions, to engage respectfully and argumentatively with interlocutors. The book forced me to think in a new way -- and to think hard with close attention! -- about a wide range of issues. In this comment I'm going to be focusing on the first part of the book, the part that is primarily about ethics.

Before reading this book, I had little previous acquaintance with the issue of internalism and externalism as those terms are understood here. Even though I consider myself a metaethicist, I tend to work on different metaethical questions than the ones found in most typical metaethics debates, and in different style that centers more on practical questions as a way in to understanding metaethical questions. I also come from a perspective that is more pluralist and less aligned with thinking about moral right answers and background truth (though of course, as Professor Weatherson discusses, the way these latter ideas relate to the book's arguments is nuanced; more on this later).

As I'll explain further below, as I was reading, I found myself in agreement with some of Professor Weatherson's claims against internalism. But I also found myself wondering what, exactly, externalism is -- or, perhaps I should say, what externalism comes to, or what this debate comes to, for the more practical questions of life. If we do not know the moral truth, how can we do the right thing, or advise others to do the right thing? If we are mistaken about the moral truth,

then we do not know the moral truth, and do not know that we are mistaken. What does normative externalism advise us to do?

Much of the ethics discussion in this book is an argument against internalism. I don't know if it's because I happened to be recently teaching about intuitionistic logic in my philosophy of math class or what, but I found myself wondering: does an argument against internalism equal an argument in favor of externalism? Is it possible that neither is right? For example, is it possible that the the questions at the center of this debate have contextually-sensitive answers, so that there is no general correctness to either externalism or internalism?

With respect to the matter of objectivity or background moral truths, Professor Weatherson says that the arguments of the book do not depend on any strong notion of objectivity and that they assume, at most, a minimal view of moral correctness that is described as conflicting with only the most extreme and implausible forms of relativism. But I was struck by how the rhetoric and framing of the discussion seems informed by a way of thinking about ethics in terms of having the right moral theory, the idea that one moral theory could be right or wrong in a certain possible world, and the possibility that one can think probabilistically about the likelihood that one's moral beliefs are, or are not, on the right track. So, for example, when we talk about someone being moral mistaken, we seem to be talking here not only about the possibility that as as a moral evaluator I or another may find their views to be wrong -- it seems also that they may just *be* wrong in some other sense. Questions of blame and judgment are brought in with a similar style. The question of whether someone is to be blamed -- or "marked down" more generally, as Professor Weatherson puts it at one point, seems here to be not only the question of whether I, or another, would find them blameworthy or worthy of being marked

down, it's the more general or abstract question that they can be blameworthy or subject to censure from some abstract point of view.

The issue of objectivity versus relativism is complex, but let me lay out for your consideration a different starting point for reflection. Professor Weatherson is perhaps correct that views in which there are no background right answers to moral questions -- that is, views in which moral truth is "relative to the moral principles of the actor being evaluated" -- face challenging objections. But I would say that most versions of moral absolutism or objectivity, and most versions of relativism in which moral truth is relative to a community, also face challenging objections. Given this, and given the role that concepts like moral truth play in normative externalism, it seems to me possible that questions about moral truth -- and what it is, and how we know it -- are not quite as disconnected from an understanding of normative externalism as is suggested. To avoid spending this whole comment on the issue of relativism versus objectivity, which one could easily do, let me set aside that way of framing the discussion, and offer instead a different perspective as a starting point. This perspective is informed by some of my previous work. It takes up a somewhat different view of moral reasoning, truth, and justification, it centers issues related to disagreement and uncertainty, and it focuses on practical questions. I'll trace out how that alternative starting point leads to a different way of seeing topics like moral uncertainty, to some of the examples related to asymmetry, and to ideas about moral community.

This alternative starting point starts from the idea that as people we value and care about certain things -- some of them people or objects, and some of them abstract values or other kinds of things. We then face the practical questions of how to reason morally, to figure out what we

think is true or justified, and how to engage with others who may disagree with us. The alternative perspective centers practical questions and the implications that metaethical views have for applied questions. As Professor Weatherson explains in the Introduction, normative externalism can be applied at the level of actions, agents, or advice. The first two are presented in abstract terms having to do with whether an action or agent is "right" or "good" in some abstract sense. My framing doesn't rule out such evaluations but it does make them a bit more challenging to understand. Right or good from whose point of view and how would we know that? The third, however, is potentially quite practical: it is whether or not we should advise people to act according to their own principles. Reading the examples and arguments in the book, it seems to me that there is another closely related practical question to the one about advice: whether we should assist someone in doing what they think is right. Blame can also be a practical question as well, and while I will touch on it some, I am going to focus more directly here on advice and assistance for reasons of space.

A shift in focus to practical questions may leads us to consider the generality of normative externalism -- and normative internalism for that matter -- in a different light. From a practical view, we might consider whether questions about advice and assistance even have general answers. What if, instead, they are context-sensitive? Perhaps the questions are whether -- and **when** -- to give one set of advice and whether -- and **when** -- to assist someone in doing what they think is right. And similarly for blame. A shift in focus to practical questions may also lead us to wonder how normative externalism can be put into practice: if we're giving advice and assistance, aren't we always giving the advice and assistance that we think is right in some sense? I return to this last point at the end.

In this context, there is special interest in cases where there is moral disagreement: if your friend thinks that some decision is morally right, and you do too, then how to advise them and whether to assist them seem to have relatively straightforward answers -- and since they've already done what you think is the right thing there is no problem about blame. In my previous work I argued that moral disagreement can sometimes be traced to differences in the way that we value. Since I want to draw on this framing to offer an interpretation of moral uncertainty, let me spend a couple of minutes explaining it here.

In my 2015 book *Moral Reasoning in a Pluralistic World*, I explore the idea that one fruitful way of understanding moral disagreement and diversity involves recognizing that the way people value is pluralistic in two senses. The first is that we value pluralistically, endorsing multiple values like benevolence, justice, liberty, and fidelity which can conflict and cause moral dilemmas. To say that values are plural is to say that there are various, genuinely distinct values, not reducible to one another or to some "super-value" (Mason 2008). It is often noted that the way we value in ordinary life seems "pluralistic" in this sense. As Thomas Nagel (1979) notes, not only do we care about various different kinds of things, but even the form of moral concerns we encounter varies. Familiar values seem to exemplify multiplicity. We care about honesty, but also about not hurting others. We care about human rights, but we also care about bringing about the collective good. In contexts of value pluralism, competing considerations can conflict: if a person must lie to keep a promise, say, then honesty may recommend promise-breaking while fidelity recommends deception. In such cases, deciding what to do is often a matter of prioritizing among conflicting concerns.

The second relevant form of pluralism is that, while we often share values, different people direct their cares at different objects and prioritize amongst them in different ways. For example, with respect to prioritization, some people may prize justice and fairness above all, overriding considerations related to the common good, while others who prioritize differently, allowing that in some cases, the collective good is most important.

In such contexts of value pluralism -- in which we recognize multiple values such as benevolence, justice, honesty, liberty, and fidelity -- I propose that some disagreement and diversity arises because while we roughly share values, we direct and prioritize those values in different ways. For an example involving prioritization, consider a case in which a person would have to lie to keep a promise. Two people may each value honesty and promise-keeping, yet disagree about what to do in this case, because they disagree over which value or obligation is most important in the given context. These people share values, but disagree over what to do in the given situation. As is often noted, the values of justice and beneficence can conflict, and we may value both while disagreeing on when, if ever, we should override the requirements of rights or respect to prevent a great harm. In the abortion debate, it is possible for people to value the potential life of a developing fetus, and also to value a woman's autonomy rights to control her body, yet completely disagree about the permissibility of abortion, because those who prioritize the latter will be pro-choice and those who prioritize the former will not. These connections between diversity and conflict help explain why so much moral disagreement seems to arise in the contexts of dilemmas, in which there is more than one value at stake, and as an individuals we feel pulled in different directions.

Working within a framework of what Geoffrey Sayre-McCord has called "conviction ethics" -- in which we start moral reasoning with our considered convictions, then aim to bring these into coherence -- I argue that in this context, for a person to bring their moral beliefs into coherence does not -- and should not -- involve norms like systematicity, in which we seek a few simple principles as possible -- but rather case consistency, which is judging in accordance with moral similarity and difference. From this framing, it follows that moral disagreements can arise in at least two different ways. Sometimes people roughly share values and prioritizations, but one person is failing to be case consistent -- that is, to judge the same way those cases they themselves would see as similar in significant ways. This can occur, for example, because of the well-known phenomenon of framing effects, where we judge differently cases that are identical in all the facts, only because of the way the cases are presented and described. But disagreement can also arise for more fundamental reasons, arising from the way people direct and prioritize values differently on a deeper level. In the latter kind of situation, disagreement can be entrenched: even when both sides are reasoning consistently and well, deep differences in value prioritization means that they will never agree. For the person who prioritizes honesty over fidelity and the one who prioritizes fidelity over honesty, and for the people who prioritize autonomy rights and the value of a developing fetus differently, coherence will not lead to consensus or agreement. In cases in which diversity in moral judgment is due to deep underlying differences in values, there can be multiply internally coherent sets of beliefs that contain different moral conclusions and judgments. In this approach we might then take the results of carrying out a coherence method as giving us a reflection of a set of values of a person or

community -- an improved version, if more coherent, but more like a reflection than an objectively correct answer.

The first thing to notice about this alternative starting point is that it suggests a potential way of understanding moral uncertainty that is different from that proposed in the book. In *Normative Externalism*, moral uncertainty is treated roughly as the idea that, in your world, there is a fact of the matter about what the moral truth is, that you yourself are uncertain what that is, and that there is some chance that you are wrong and some chance that you are right. Occasionally a larger frame is brought in, in which moral uncertainty is intimately linked to not knowing which comprehensive moral theory is right -- for example, should we be utilitarians or Kantians? From my perspective though, moral uncertainty can arise in a different way, of being unsure how to prioritize different values, or seeing them as close to evenly matched, even though one is committed to both.

Professor Weatherson draws on moral uncertainty in his argument against the symmetry thesis -- a thesis which is often used to support internalism and which draws a parallel between moral uncertainty and factual uncertainty. Professor Weatherson has several arguments against the symmetry thesis, some of which focus on examples like the abortion example. In this example, Marilou lives in a state where it is illegal and extremely difficult to get an abortion and she wants to have one, and she asks her friend Shila for assistance. Shila is morally uncertain about the moral status of abortion at Marilou's stage of 12 weeks of pregnancy, and is unsure whether it is a permissible act or whether it is the moral equivalent of murder. Shila can either assist Marilou, report Marilou to the authorities (which will likely lead to Marilou's imprisonment), or do nothing -- in which case she expects that Marilou will give birth and give

the child up for adoption. Professor Weatherson argues that if there is symmetry between factual and normative uncertainty, then the last option is the right one. In factual uncertainty, we should often hedge our bets and avoid doing something catastrophic by doing something safe but non-optimizing. If symmetry were right, we should hedge our moral bets and do nothing. But Professor Weatherson argues that hedging our bets and doing nothing in this case would just be wrong: it would be cowardly and would not be a safe or prudent choice. Shila, Professor Weatherson says, should either help Marilou or report her. This undermines symmetry and thus undermines a main argument in favor of internalism.

From the alternative framing above and from my own point of view, the abortion example lends itself to a different interpretation -- one that does not support symmetry, but for different reasons, and one that does not seem to me to support externalism either. One possibility -- though of course not the only one, an interesting one -- is the case in which Shila values both the developing fetus and also Marilou's autonomy to make decisions about her own body. She sees the issue in light of a moral conflict, in which two values cannot both be fully honored. Shila is unsure, at this time, how to weigh these values against one another in the given case. To make the case vivid, we can imagine that Shila is in the midst of a values-shift in her life. Perhaps she has grown up with one perspective on the abortion issue in general and has only recently started to think about alternatives. If she sees the values in the case as close to evenly matched, she may be able to sympathize in a meaningful way with both people who think the abortion would be wrong and also those who think that it would be perfectly acceptable.

Perhaps we are now getting to just conflicting intuitions, but my sense is that if I were Shila, I might go in one of a few directions, all diverging from the analysis suggested by

normative externalism. First, I might think that as someone in the midst of changing my mind, who is not clear on this issue, I would be the wrong person to be a deciding factor either way. So I might think that this is indeed a reason to hedge and stay out of it. Second, and more likely, I might ask someone else for advice. That is, I would see the occasion as one for moral humility, and consult with someone else whose values I trust. Third, and most likely still, I might choose to regard the bonds of friendship as an overriding factor in the case, and do as Marilou asks. This third option may involve seeing the case as one in which values are evenly matched, so that other considerations become a determining factor.

Does my way of seeing things support the symmetry thesis and thus internalism? No, in that I think the option of taking friendship as a deciding consideration means helping Marilou instead of staying out of it, and this fits with Professor Weatherson's conclusions against internalism. But: does my way of seeing things fit with normative externalism? I think the answer to that question is no as well. I am not sure exactly what normative externalism tells us in this case, but I suppose it is something like this: if abortion at Marilou's stage of pregnancy is morally permissible, Shila should help; abortion at Marilou's stage of pregnancy is wrong, Shila should report her friend. If Shila is as I have described her above, then she either stays out of it and does nothing (obviously not externalist) or she helps for reasons that context-sensitive and have little to do with moral truth -- reasons like the advice she got, or the bonds of friendship acting as decisive factors in moments of moral uncertainty. This analysis seems in tension with the one normative externalism would suggest.

This example suggests that cases of moral uncertainty that arise because a person sees a situation as a conflict of values and either 1) sees them as closely balanced with considerations

on both sides or 2) isn't sure how to balance them are cases in which our decisions having to do with advice and assistance are going to be especially nuanced. They also seem like cases in which moral hedging or caution may be appropriate considerations. These reflections don't support internalism, as they advise hedging and caution on other grounds than the symmetry thesis, but they don't seem to support externalism either, as the reasons for hedging or caution do relate directly to uncertainty. If the answers to questions about advice and assistance are context-sensitive, is it possible that neither internalism nor externalism is correct?

Perhaps from a practical point of view, normative externalism could be interpreted as the idea that when giving advice and assistance, we should not say to a person "you have to do what you think is right." As Professor Weatherson focuses more on agents and actions than advice, I'm not sure what the alternative is and whether the correct alternative is to say what *we* think would be right -- as I mention briefly below, there is a sense in which this alternative seems internalist as well, as we would be advising the person according to our own sense of what is best to do. But let me set that aside here and consider the possibility that normative externalism tells us that advice should be given from the moral point of view of the advisor and should not point the advisee toward doing what they think is right.

With the question put this way, from my perspective, I would say that questions of advice and assistance may vary from case to case according to a wide range of particulars. Certainly, there are many cases in which I would advise someone to do just as *I* think they ought to do in a case, and would not say "you have to do what you think is right." A friend asked me for help cheating her insurance company; another asked if he should misrepresent a fact in job application; a third wondered whether they should give a student an inflated grade to help them

out (these are all real examples). I told them all "no." I am not sure whether this is in line with normative externalism -- I was of course doing as I thought right, myself. But perhaps it is consistent with not-normative internalism: I did not give them the advice "you have to do what you think is right." I gave them the advice to do the right thing -- or, what I took to be the right thing, anyway.

But other examples seem to me more complex. A friend asked me whether she should reconnect with a parent who had been abusive and who was longing to speak with her; another faced a decision about whether to pursue further treatment for a family member who had had a stroke and would be unlikely to regain the ability to move, talk, or eat on their own; a third had to figure out how to engage respectfully with their religious Muslim parents as an agnostic and proponent of gender equality. A friend of a friend was trying to work out whether to bring charges against an older person who had had sex with them when they were legally too young to consent -- they had regarded it at the time in light of a consensual interaction, should they now initiate a process in which that person might go to prison? (These are also all real cases). In these cases it seems to me -- to me personally, as a moral evaluator with my own values and beliefs -- that there is a lot to be said on various sides, and that it seems appropriate to say to a person that they should do what they think is right. Further, I would be willing to assist them in doing so -- even if they see the values as weighted different from the way I see them.

My sense is that there are times when the right advice is that someone should stick to their guns and there are times when the right advice is to have a little more moral humility. These are questions about how we relate to one another in difficult circumstances and they seem to me to hinge on a wide range of complex contextual factors. This is purely observational, but

descriptively my sense is that many people's moral views have subtle distinctions embedded in them regarding what we might call "the moral penumbra" -- the zones in which we have beliefs but accommodate difference as contrasted with the zones in which we have belief and cannot accommodate difference. If that way of doing things has something to say for it, then maybe there is no generalizable answer to the question of what kind of advice makes sense.

In the chapter on blame and whether ignorance is exculpatory, Professor Weatherson comes closest to these kinds of thoughts when discussing people in our moral community versus people not in our moral community. There, he suggests that externalism is general is committed to the view that moral ignorance is not exculpatory. But here he discusses JoJo -- a vicious dictator who was raised by a dictator to know of no other way of life. A widely shared intuition, Professor Weatherson says, is that JoJo is less blameworthy than his father, a dictator who did know another way of life who was installed after a coup. Professor Weatherson says that one potential explanation for this intuition -- one way of understanding it -- is that JoJo is outside our moral community, and thus we do not judge him the same way we judge the people in our moral community. This has echoes of my examples, which also involve people who value differently from the way we do -- though I am not sure from the context whether Professor Weatherson means to include people who value differently from us as people who are outside our moral community or whether it is just cases like JoJo who face more complex situations. In any case, I was struck that Professor Weatherson uses the phrase "moral strangers," which is also used by Tristram Engelhardt in the bioethics context for people who value differently. I -- and Engelhardt -- mean to refer both to people outside our social community who value differently from us, and also to people inside our social community who value differently from us. For my part, I've often

thought that the phrase "moral strangers" (contrasted by Engelhardt with "moral friends"), may overstate the differences between us and people with whom we share values but disagree on prioritization. But the main question here concerns how normative externalism handles blame relating to people who value differently -- whether that counts as being outside of our moral community, and whether we change our attitudes with respect to blame in such cases.

As mentioned briefly above, focusing more on the practical questions of advice I also wondered about how giving advice looks from the externalist and internalist perspective, because there seems to be one sense in which every way of proceeding seems, in some sense, internalist. In the examples I've given, I am of course drawing on my own sense of when it is right to give advice one way and when it is right to give advice another -- and even if I took someone else's advice about *that*, I would still be deciding to do so. So -- setting aside the possibility that weakness of will somehow causes me to give advice that even I don't think is right -- once the question is framed in a more practical way, my sense of what the question or debate is starts to become less clear. Is internalism being interpreted instead as saying that we have some extra reason to do what we think is right, beyond the fact that it is what we will have to end up doing in some sense? If yes, that seems to me to put a different kind of pressure on the question of whether not-Internalism is really the same as externalism: not emphasizing the requirement to do as you think best does not seem to equal a requirement to do something else.

In summary, I've considered an alternative starting point for reflection that has a different perspective on moral truth, I've suggested that it leads to a different way of seeing moral uncertainty in which answers to questions about advice and assistance may not generalize, and discussed the idea that in that case, perhaps neither -ism is correct. I've also briefly raised the

possibility that seeing advice in practical terms (and setting aside weakness of will), we're already in some sense doing what we think is right: is internalism being interpreted as the idea that we have extra reason to do so?

In any case, I am grateful to Professor Weatherson for laying out these issues in such a crisp clear way, for the illuminating back and forth with philosophers he disagrees with, and for a book that is thought-provoking in the best way possible.