(This is a draft from 2014.)
Published in *The Philosopher's Magazine* (66): 94–99, 2014.

# Not Knowing Everything That Matters

Jonathan Dancy and Daniel Muñoz

There's more than one way to get things morally wrong. If I take candy from a baby and throw it (the candy) into the dirt for my own amusement, knowing full well what I'm up to, I get things wrong in at least two ways. My action is a display of excessively selfish character, and it leads to an unfortunate result. Now consider a different case, in which my beliefs and body movements are the same, but I don't know a relevant fact: the candy is laced with a deadly poison. Here, I still deserve to be criticized as selfish, but it's no longer true that I've harmed anyone—on the contrary, I've saved an innocent child's life! In another case, I'm no longer selfish, so I leave the baby to enjoy its treat, but the candy is still poisonous. At least I'm displaying good character, but due to my ignorance, this noble display leads to a tragic result. And in yet other cases, I get things wrong by ignoring my evidence—perhaps I notice that the candy is marked 'POISONOUS TO INFANTS', yet fail to infer that the baby is in danger—or I err because of my sincere but twisted moral beliefs, such as the belief that one ought never to be kind to anybody under the age of 11.

When I don't know all of the relevant facts about my situation, when I fail to use my evidence, and when I hold false moral beliefs, my ignorance of what matters leads me to get things wrong in distinct ways, and I come to deserve distinct kinds of moral criticism. How are we to understand these various failures and censures?

In Ch. 7 of his monumental *On What Matters*, Derek Parfit proposes an answer: for each way of getting things wrong, there is a different *sense* of the word 'wrong'. He begins the chapter by saying that "in many cases we don't know all the relevant facts, and we must act in ignorance, or with false beliefs. When we think about such cases, we can use 'wrong' in several partly different senses. Some of these senses we can define by using the ordinary sense. Some act of ours would be

> *wrong* in the *fact-relative* sense just when this act would be wrong in the ordinary sense if we knew all of the morally relevant facts,
>
> *wrong* in the *belief-relative* sense just when this act would be wrong in the ordinary sense if our beliefs about these facts were true,

and

> *wrong* in the *evidence-relative sense* just when this act would be wrong in the ordinary sense if we believed what the available evidence gives us decisive reasons to believe, and these beliefs were true. …

We ought to use 'wrong' in all these senses. If we don't draw these distinctions, or we use only some of these senses, we shall fail to recognize some important truths, and we and others may needlessly disagree. …" (pp. 150-1)

Later Parfit adds a fourth sense of wrong.

"Some act is

> *Wrong* in the *moral-belief-relative* sense just when the agent believes this act to be wrong in the ordinary sense." (p. 158)

These four definitions of *wrong* all appeal to what Parfit calls "wrong in the ordinary sense". Unfortunately he doesn't tell us what that sense is. He does mention an 'indefinable' sense of *wrong* that can be explicated by appeal to the notion of what 'mustn't be done'. (p. 165) Perhaps, then, it does no harm to read 'is wrong in the ordinary sense' as meaning 'mustn't be done'.

Our first question should be whether we are really dealing with different senses of 'wrong' here. Is it not simply that each of these different senses is really definable in terms of one and the same sense? If so, that sense would of course be Parfit's *fact-relative* sense. But this pleasing prospect has to be abandoned. In each case, we are dealing with a sense of 'wrong' under which it is true that our act is wrong in that sense if it would (under certain conditions which do not obtain) be wrong in the *fact-relative* sense. This does not merely mean that there would be something conclusive to be said against doing it if those conditions obtained – which they don't. It means that there is in fact something to be said against doing it (so it is 'wrong in a sense') just because if things were as the agent supposes (in non-moral respects) there would be decisive reason not to do it – it would be something that mustn't be done. This 'something to be said against doing it' is there even though in fact the action is not wrong in the ordinary sense at all.

But should we really agree with Parfit that there is a sense in which someone's action is wrong if it would be wrong in the *fact-relative* sense under certain conditions that do not obtain? Yes we are, because this is the only way to retain the sense that there is something criticisable about an action done by someone whose beliefs are such that, if they were true, she would be acting wrongly, even though those beliefs are not all true and she is not acting wrongly. (Recall the case in which I believe I'm stealing candy from a baby, when in fact I'm preventing it from eating poison.) If we simply say that she is not acting wrongly, it sounds as if we are in no position to criticise her for so acting, unless we can find some other derogatory epithet with which to

stigmatise her. So we seem to have two possible resolutions. One is to say is that *in another sense of 'wrong'*, she is acting wrongly, even though in the *fact-relative* sense she is not. This is Parfit's resolution. The other resolution is to find a range of derogatory epithets, one for each sort of case. If we take this route, we might reserve 'wrong' and 'ought not' for the fact-relative case and say of the other cases that, though the act in those cases is not wrong, it is still subject to some form of criticism, which may be more or less severe, ranging from outrageous and unjustified through reprehensible to inappropriate. This second option is likely to be fairly stipulative. There probably is not a nice set of critical epithets perfectly fitted to cover the ground. We would have to invent one. So this second option does not look very promising.

Returning then to Parfit's way of doing things: one might wonder whether we are really driven in this way to find a different sense of 'wrong' for each sort of case where we find something to disapprove of. Is there no other approach that will deliver the goods required? One such possibility makes do with a single sense of 'wrong', but relativises the use of that sense to different things in different cases. So there are various questions that we can ask using a single sense, of which these are some of the most important:

Would it be wrong for her to do this, given all the facts?

Would it be wrong for her to do this, given what she believes?

Would it be wrong for her to do this, given the evidence available?

Would it be wrong for her to do this, given the evidence available to her?

Would it be wrong for her to do this, given her beliefs about what evidence is available?

Would it be wrong for her to do this, given what we know (even if she doesn't)?

Of these questions only the first uses an unrestricted conception of what is 'given'. The others all use a restricted conception of what it would be wrong of us to do, 'given only' this or that more or less limited conception of what we might call the base. So all but the first use a 'given only' operator, and this operator seems to make some sense. And if it does make sense, then we can express all the questions we want to raise without multiplying senses of 'wrong'.

The same structure is to be found if we consider the sense of 'ought'. There are just various questions we can ask, among which are:

What ought she to do, given all the facts?

What ought she to do, given only the evidence available?

> What ought she to do, given only the evidence available to her?
>
> What ought she to do, given only her non-moral beliefs?
>
> What ought she to do, given only her beliefs about what evidence is available?
>
> What ought she to do, given only what she knows?

The first of these expresses Parfit's *fact-relative* sense of 'ought. He writes later that

"We can call some possible act

> what we *ought practically* to do in the *fact-relative* sense just when and because this act is what we have decisive reasons, or most reason, to do."

And he proceeds to lay out other supposed senses of 'ought', analogous to the various senses of 'wrong' that we have seen above.

Does this appeal to 'given only' clauses succeed in making sense of the relevant aspects of our evaluative practice, especially our desire to disapprove of people on occasion for doing actions that are not wrong in the *fact-relative* sense, or which we cannot say that they ought not to have done, again in the *fact-relative* sense? We don't see why not. We can surely disapprove of someone who does what she ought not to have done if all her beliefs had been true, that is, if the situation had been in every respect as she conceived it to be. And the reason why we disapprove of her is exactly this: that given only the world as she conceives it, her action would be wrong. We can also retain a way of disapproving of someone who does what is in fact right, but which would have been wrong if things had been as she supposed. And we can find something nice to say about someone who does what is in fact wrong but which would have been right if things had been as she perfectly reasonably took them to be. What is more, we don't need to find a whole raft of epithets, one for each case. We make do with one concept, the concept of 'what is right given ….', which enables us to distinguish all the cases just as we would wish.

What we see here, then, is that the need to retain a way of criticising people who do actions that are in fact exactly what they ought to do, given all the facts, is quite easily met. We need only to hold that what we mean by 'wrong' is sensitive to a certain feature of the context in which we use it—namely, whether the considerations 'given' are the facts, someone's beliefs, someone's evidence, etc. So should this not be enough to persuade us that it is better to have a single concept of 'wrong' and a single concept of 'ought', so long as we can use these concepts to ask all the questions we want to ask? Better, that is, than having a list of such concepts so that according to some the action was wrong and according to others it was not and might even have been what he ought to have done? Alas, matters are not so simple.

The reason why they are not so simple is that the phrase ‘given only’ is very hard to make good sense of. What could it mean to say that, given only what he knew, his action would have been wrong? It seems to require that an action can be made wrong by a limited group of facts, irrespective of how things are in other respects. Let us examine this point in more detail.

Consider the sentence ‘given only that she promised to do it, she ought to do it’. This seems to mean that her having promised has a sort of ought-making capacity. This doesn’t mean that one ought always to do what one promised to do; one ought not to keep one’s promise if doing so would lead to great disaster. But it does mean, or at least looks as if it means, that if there were no other facts at all beyond the fact that she promised, she ought to do what she promised to do. But we cannot make sense of the idea that there are no other facts at all. A world that consists of only one fact is inconceivable. So what does the given-only sentence mean? It might mean that if she has promised to do it and the world contains no other facts that make any difference to what she ought to do, she ought to keep his promise. One trouble with this is that it is not obvious why it is even relevant to what she ought to do in the world as it stands, since the actual world no doubt contains lots of things that make a difference to what she ought to do. But another trouble is that it is impossible for a world to contain no other facts that make any difference to what she ought to do. A promise is made either freely or under constraint. If it is made under constraint, this affects the question whether the promiser ought to do what was promised. If it was not made under constraint, that fact is also relevant. Does it make a difference to what she ought to do? Yes it does, since if it had not been so, a difference would have been made. So there is no such thing as a world that contains no other facts that make any difference to what she ought to do, and so there is no such thing as the difference made by the mere fact that he promised. And presumably this argument is entirely general. Every ought-making fact is such that under certain circumstances it would not have been an ought-maker. If that is so, there can be no answer to the question what one ought to do given only that fact.

Try another example, just to check the point. What ought you to do, given only that there is a child drowning in the pond in front of you? What we want to say, of course, is that you ought to jump in and save her. But this is not because, if the world consisted of only one fact, and that fact was that there is a child drowning in the pond before you, you ought to save her. Nor is it because, if the world consisted of the fact that there is a child drowning in the pond before you plus the fact that there are no other relevant facts, you ought to jump in. There is no such thing as the fact that there are no other relevant facts. There will always be a further relevant fact. For suppose that, though there is a child drowning in the pond before you, there are also two children drowning in the pond just behind you (or simply, more generally, that you have another more pressing and competing duty). That would be a relevant fact. If so, the fact that you have no more pressing and

competing duty is a relevant fact about (or in) the first world. So there will always be some further relevant fact.

The same point also thwarts another slightly different way of making sense of ‘given only’ operators. Suppose we interpret ‘given only that she promised to do it, she ought to do it’ to mean ‘in a world where she promised to do it (ignoring the moral relevance of other facts), she ought to do it’. Crucially, the moral relevance of the promise itself—whether it binds her to keep it—will depend on whether the promise was made freely or under duress. So, even though we’re only considering the relevance of the fact that she promised, if we ignore facts about the conditions under which the promise was made, we will be unable to determine whether she ought to keep her word. There will just be no fact of the matter.

The conclusion seems to be that the very idea of a ‘given only’ sentence makes no sense. And if so, we are back in the original unsatisfactory position of having to choose between (a) Parfit’s list of many senses of *wrong* (and of *ought*), and (b) finding a list of appropriate epithets to characterise the different ways in which people may be criticisable for doing what was in fact the right thing to do, or praiseworthy for doing what was in fact the wrong thing to do. This conclusion is not one with which theorists should be at all happy.

Jonathan Dancy teaches philosophy at the University of Texas at Austin and is Research Professor at the University of Reading, UK.

Daniel Muñoz was an undergraduate at UT Austin and is about to start as a graduate student of philosophy at MIT.