

Introducing Metaethics

Zoë Johnson King*

Department of Philosophy, Harvard University, Cambridge, USA

*Corresponding author. Email: zjohnsonking@fas.harvard.edu

Keywords: metaethics; relativism; morality; God; wrong

Abstract

We often describe actions as *good*, *bad*, *right*, *wrong*, *fair*, *unkind*, *deserved*, *disrespectful*, *a bit much*, and so on. This article asks: Do these terms describe facts about our actions? And do those facts tell us to perform certain actions and refrain from performing others? If so, what exactly does that mean? And, if not, what are we doing when we describe actions in these various ways?

It's strange to think that there might be facts that tell you what to do.

Facts normally seem to be *just sitting there*, out in the world. The fact that grass is green doesn't tell you to do anything. The fact that penguins primarily live in the Southern hemisphere doesn't tell you to do anything. Facts about what colour things are and about who lives where are just out there minding their own business.

How about facts about what is *right* and what is *wrong*? Suppose that self-isolating when you test positive for a highly infectious disease is *the right thing to do*. Does this fact about rightness tell you to self-isolate when you test positive for a highly infectious disease? Likewise, suppose that storming the US Capitol when you're angry because your candidate lost an election is *wrong*. Does a fact about something's being wrong tell you not to do that thing? This is a puzzling question. On the one hand, it seems to be part of the meaning of the words 'right' and 'wrong' that you have to do what's right and refrain from doing what's wrong. On the other, it remains difficult to wrap one's head around the idea of a fact telling you what to do. Aren't

these facts about what's right and what's wrong also just sitting there for us to respond to as we wish, like the facts that grass is green and that penguins primarily live in the Southern hemisphere?

Actually, come to think of it, *are* these facts just sitting there? Where are they sitting?

Facts about rightness and wrongness are not the sort of facts that we detect using the five senses or any of the usual methods of the physical sciences. So it's unclear whether these facts are part of the physical world at all. It's unclear whether they're sitting anywhere, so to speak. And, in that case, one might worry that these aren't really facts at all. Indeed, one might worry that one doesn't really know what a fact *is*. At first blush, it seems as though a fact is just a state of affairs corresponding to a thought that one might entertain and making it the case that this thought is correct or true. But surely it *is* true that self-isolating when you test positive for a highly infectious disease is the right thing to do. So where's the fact?

Now we're doing metaethics. Metaethics is the best and coolest subfield within Philosophy. It's about what we're up to when we say things like 'metaethics is the best and coolest subfield within



Philosophy'. Metaethicists ask: when we make claims about what's *good* or *bad* or *better* or *best* or *worse* or *worst*, or about what's *right* or *wrong* or *okay* or *pretty darn nice* or *going above and beyond*, or about what's *appropriate* or *deserved* or *called for* or *warranted*, or about what's *cool* or *lame* or *beautiful* or *striking* or *dull* or *boring*, or about what's *fair* or *just* or *over-the-top* or *belittling* or *respectful* or *sanctimonious* or *a bit much*, what are we talking about? Are there facts that make what we say true (or false)? If so, what kind of facts are those, and how do we detect them? If not, what are we doing when we make all these claims?

OK, so let's go back to the whole facts-telling-you-what-to-do business. You might think that the problem was that we were looking at the wrong facts. Maybe it's not the facts *that* certain actions are right and others are wrong that tell us what to do, but some other facts that explain *why* the right actions are right and the wrong actions are wrong. Maybe the rightness-facts

and wrongness-facts are more like verdicts that point out that we've been told to do some things and avoid others, with other facts doing the actual telling.

But the trouble is that even the facts that seem to explain why certain actions are right and others are wrong don't seem to be able to tell us what to do. For example, maybe storming the US Capitol when you're angry because your candidate lost is wrong because it is *undemocratic* and *unjust* and *unfair*, or something along these lines. And maybe self-isolating when you test positive for a highly infectious disease is the right thing to do because it's *considerate* and is *your way of doing your part*, or maybe it's just *the least you can do*. These facts about what is unfair, what is the least you can do, and so on aren't facts that it's easier to find in the world than facts about rightness and wrongness. So that's not much help. Alternatively, maybe self-isolating when you test positive for a highly infectious disease is the right thing to do because it

slows the spread of the disease and thereby protects others around you from infection. This seems like a respectable, scientific sort of fact. So that's progress. But how can these scientific facts tell us what to do? Aren't they just sitting there like the fact about penguins, to be responded-to as we wish?

'Facts about rightness and wrongness are not the sort of facts that we detect using the five senses or any of the usual methods of the physical sciences. So it's unclear whether these facts are part of the physical world at all. It's unclear whether they're sitting anywhere, so to speak.'

At this point people sometimes appeal to God. But it's not obvious that the fact that God commanded you to do so-and-so fares better than any other (putative) fact about what the world is like and what has happened in it. The fact that God commanded you to do so-and-so is a fact about what someone has said, like the fact that your crush invited you to see a film or the fact that your little brother asked you whether dogs can skateboard. These facts about things people said might make you *feel* like doing something, especially if you are favourably disposed towards the people who said the things. But do the facts *themselves* tell you what to do? Is the

fact, itself, a command? That doesn't seem right. The fact *is* that God commanded you to do so-and-so. So it isn't the fact that commanded you – it's God!

Why should you do what God says, though? God might just issue commands willy-nilly, commanding whatever they feel like commanding regardless of whether it's genuinely a good idea. But then the fact that God commanded you to do something is like the fact that your school says you can only wear a bobble if it's green; that's something that you might take into account in your actions just because you don't want to get detention, but it doesn't imbue the eschewing of non-green bobbles with any broader cosmic significance. School uniform colours have an obvious arbitrariness to them. And if God's commands are similarly arbitrary, then they seem similarly devoid of cosmic significance, even though you might follow them anyway just because you want to avoid eternal torment. What's more, if you decide that you actually don't mind detention or eternal torment, then you can just go ahead and violate your school or God's commands. So these facts about commands seem to be paradigm cases of facts that don't compel us to do anything in particular and that we may respond to as we wish.

Maybe God is all-knowing and all-loving, so they would only tell you to do something if it genuinely was a good idea. Then the fact that God commanded something doesn't have the same arbitrariness as facts about school uniform colours. It's more like the fact that your school says you can't bite people's legs while they're getting on with their work – here the school is telling you do something that is independently a good idea. But if God only tells you to do what's independently a good idea, and God tells you to self-isolate when you test positive for a highly infectious disease, then self-isolating when you test positive must be an independently good idea. So, where's that independent fact? The fact that self-isolating is genuinely a good idea can't consist in God's command if God is *responding* to that fact when he commands you. And now we're right back where we started, looking in the world for facts about goodness without really knowing what to look for. Introducing God didn't help.

Another classic move at this point is to go *relativist*. We might be tempted to say that there are no ‘objective’ facts about what is right, wrong, pretty nice, fair, respectful, sanctimonious, and so on, but only facts about what is right (etc.) ‘for’ particular individuals – or for societies, families, cultural traditions, or other groups. We might be tempted to say, then, that self-isolating when you test positive for a highly infectious disease is the right thing to do *for us*, given our values and beliefs, and is not the right thing to do *for* another individual or group who rejects our values and beliefs. But this view is in danger of collapsing in on itself. If there are no objective facts, then what is the status of the relativist’s claim that certain actions are right for certain people and wrong for others? This claim is a claim about what’s right. So the relativist cannot coherently say both that this claim states an objective fact and that there are no objective facts about what’s right. That view is internally inconsistent; one half of the view entails that the other half is false.

The alternative is to be a relativist about relativism itself. We could say that relativism is true *for* relativists, given their values and beliefs, whereas there are objective facts *for* people who believe in such things. But this view is hard to wrap one’s head around. It’s hard to know how there can be objective facts ‘for’ some people but not for others; objective facts are meant to be precisely the sort of thing that don’t pop in and out of existence depending on who we’re looking at. Moreover, relativism about relativism seems to give people far too much credit. It suggests that we can *change the facts* about metaethics (for us) just by changing our metaethical beliefs. We’re not that powerful. Surely some people just have metaethical beliefs that are incorrect. That’s why we worry about metaethics; we worry about whether we’re correct or incorrect. If you lay awake at night worrying about whether anything is really wrong, I doubt that this worry could be assuaged by the cheery thought that if you simply choose to believe in objective moral facts then – poof! – there will be some *for you* (though not for the relativists).

A version of this problem applies to regular old relativism about rightness and wrongness and

a-bit-much-ness and so on, just as it does to relativism about relativism. For this view also seems to give people too much credit. After all, some pretty wacky sets of values and beliefs have appeared over the course of human history. And even wackier ones are possible. It is hard to believe that all of these values and beliefs reflect ‘facts’ that are true for the people who adopt them. For example, suppose that someone thinks that the only thing that matters is that we maximize the number of porcupines in the world. Surely it gives her too much credit to say that porcupine-maximizing *is* the right thing for her to do. She *thinks* it’s right, of course. But isn’t it plausible that she’s simply incorrect? If even God’s commands can’t imbue stuff with cosmic significance then it is hard to see how this fool could do better than God. It’s hard to see how we could change facts about metaethics just by changing our metaethical beliefs, and similarly hard to see how we could create facts about ethics just by having some ethical beliefs.

Perhaps this means that it was a mistake for the relativist to talk about facts at all. She could stick with beliefs instead: rather than saying that actions are right *for* certain people and wrong *for* others, she could say only that they are right *according to* those people and wrong *according to* those others. But this much is uncontroversial. Everyone agrees that people have moral beliefs. This is one of those facts that doesn’t tell you to do anything, like facts about grass and penguins. When something is wrong according to your values and beliefs you might *feel* like avoiding it, much as when your little brother has asked you whether dogs can skateboard you might feel like answering him. But the facts themselves don’t command or compel you to do these things. You just do them because you feel like doing them. If you were in a spiteful mood, you might instead recognize that your little brother has asked you whether dogs can skateboard and feel like totally blanking him. Similarly, you might recognize that something is wrong according to your values and beliefs but still *really* feel like doing it sometimes. Once again, then, the facts about what is right (etc.) *according to us* seem to be facts that we may respond to as we wish.

**‘Now we’re doing
metaethics.
Metaethics is the best
and coolest subfield
within Philosophy. It’s
about what we’re up
to when we say things
like “metaethics is the
best and coolest
subfield within
Philosophy”.’**

One might try denying that it is possible to recognize that something is wrong according to your values and beliefs but still really feel like doing it sometimes. After all, believing that something is wrong normally makes you strongly averse to doing it. This might be because part of what it is to believe that something is wrong is to be strongly averse to doing it. And, in that case, your feeling like doing it just goes to show that you don’t fully believe that it’s wrong – you have a pseudobelief, or at best a partial belief. This would mean that beliefs about what is wrong are *very* different from beliefs about what is green. For beliefs about what is green do not consist, even partially, in motivational states; you can fully believe that something is green and be left completely cold by that fact. The belief that something is green is just a representation of the world. So perhaps what is distinctive of ‘beliefs’ about wrongness, badness, a-bit-muchness, fairness, deservedness, respectfulness, and so on is that they are *not* just representations of the world but also something else; something like an emotion or a commitment to acting in a certain way, which brings motivation in its train.

Some metaethicists actually think that ‘beliefs’ about wrongness, badness, a-bit-muchness, fairness, deservedness, respectfulness, and so on

consist *wholly* in emotions or commitments or other mental states that bring motivation in their train. On this view, there are no facts that tell you what to do. Instead, the mental states that looked like representations of these facts – mental states that we express by describing things as right, good, unfair, respectful, a bit much, deserved, kind, thoughtful, impatient, and so on – are mental states that *themselves* get you to do things. We *express* these motivational mental states using language that makes it look as though we’re describing properties of the world outside our heads. But we’re not really. Our language just indicates that we ourselves feel a certain way about certain parts of this world and we stand ready to do something about it. The facts don’t tell us what to do. But that’s okay, because we want to do something anyway.

This view easily explains why believing that something is wrong normally makes you strongly averse to doing it. But it has a hard time explaining why the word ‘normally’ appears in that last sentence. If the ‘belief’ that something is wrong *just is* a mental state that brings motivation in its train, then why does it only *normally* make you averse to performing the action? Why not *always*? For example, consider someone who says they’re convinced that eating meat is wrong but they do it anyway because they love the taste. Or consider someone who says they know that carrying around their empty crisp packet until they find a bin would be the conscientious thing to do, but they can’t be bothered, so they just drop it on the side of the road. These people look as though they have beliefs about what is wrong and what is conscientious that bring no motivation in their train. It is possible that they don’t really have the beliefs, of course; they could be self-deceived or lying. But isn’t it also possible that they believe what they say they believe and they just aren’t motivated to act accordingly because they’re selfish and lazy? If such a person is possible then beliefs about what is wrong and what is conscientious (and so forth) cannot consist *wholly* in motivational mental states. Sometimes they are just plain old representations of facts about the world – representations that may fail to stir anything within us.

Indeed, beliefs about wrongness, goodness, conscientiousness, appropriateness, sanctimoniousness and their ilk behave like plain old representational states quite often. For we often give sophisticated arguments for the conclusions that certain actions, outcomes or people possess these various properties. And we engage in complex reasoning about which of these conclusions are true and which are false, comparing considerations that support them with considerations that challenge them. We think, for example, about whether it is *appropriate* to spend the same amount of money on presents for our two cousins even though one is older than the other; about whether it is *wrong* to keep a library book after its due date if you know that nobody else has requested the book; about whether it is *fair* for people who live below the poverty line to be punished for stealing basic necessities that they cannot afford to buy. It is difficult to make

sense of this reasoning if we maintain that the ‘beliefs’ involved are all just emotions or commitments. Instead, they look very much like representations of things’ properties and of the further properties that explain why they have those properties – much like representations of grass’s greenness and the surface reflectance properties that explain why grass is green.

Ultimately, it might turn out that facts about what is right and what is wrong (and their ilk) just *don’t* compel anyone to do anything. There might be nothing special about rightness and wrongness that makes us care about them. Maybe most of us are good people who do in fact care about doing what’s right and refraining from doing what’s wrong, and that’s all there is to it. In that case, we wouldn’t be making any sort of mistake if we simply gave up these concerns. So I suppose we should be very grateful that we, and those around us, do in fact have them.

Zoë Johnson King

Zoë Johnson King is Assistant Professor of Philosophy, Harvard University.