

I. First Reflections on Cases and on the Nature of Ethics

We're interested here in metaethics: the study of the nature of ethics; and in moral psychology, the issue of what is happening when we think about ethics. But we'll start with some issues in normative ethics.

A. Some cases

The trolley cases (Foot, Thomson, many—too many—others)

First case: the trolley is on track to kill five workers; it could be diverted onto a side-track where it would only kill one. Are you permitted to divert it? Are you required to?

Second case: the trolley is on track to kill five workers; it could be stopped by throwing one (large) worker into its path. Are you permitted to throw this worker? Are you required to? Would it make a difference if instead of a worker you threw a bystander?

Third case: like the first, except that the side-track loops round and rejoins the main track. If the trolley were to continue round, it would go on to kill the five workers. Suppose first there is a rock on the track that will stop it. What should you do? Now suppose that there is no rock, but a worker who is large enough to stop it. Are you still permitted to divert the trolley? Are you required to? Could you induce the worker onto the side track if he weren't already on it? Would you be required not to warn him if he were already there? Could you force him back onto the track?

Fourth case: Like the first, except that there is a second side track on which *you* are located. So if you leave things to run their course, five will die, but you could divert it either so that it killed you, or so that it killed someone else. Are you permitted to divert the trolley so that it kills the other person rather than you? Are you required to divert it onto yours? What lessons should we draw from this about the first case?

Some possible factors influencing our judgements: a difference between doing and allowing; a difference between treating someone as a means and treating them as an end; a difference between intending an outcome, and foreseeing it as an unintended (and perhaps unwelcome) consequence of one's actions; a prohibition on treating people in certain ways (on violating their rights); a difference between diverting an existing threat, and creating a new one.

B. The nature of our intuitions

Our intuitions seem to be of two kinds. We have reactions to *particular cases* (you should divert the trolley in the first case; you shouldn't push the person in front of it in the second); and we have reactions to different *moral principles* (it's the intention that counts; doing is more significant than allowing; you shouldn't use people as means). How do these fit together? How should we react when they clash? How much is the process like constructing a scientific theory? Should we expect to be able to construct a normative moral theory?

Stepping back a bit: what is an intuition here? Psychologists tend to mean something like a quick gut response (cf. Type I states). Perhaps it is driven primarily not by a belief but by an emotional state. Philosophers, in contrast, tend to mean something that is epistemically fundamental: an axiom or suchlike. So these may have been the result of a great deal of careful thought.

Are our intuitions *sources* of moral knowledge? Or are they attempts to *codify* our responses? If the latter, what grounds do we have for expecting that they will be amenable to systematic codification?

C. The nature of moral theories: two possible views

Ethics doesn't look to be *empirical*; it's hard to think how to do experiments to test it. But does that rule it out as knowledge? Compare mathematics. Maths is *a priori*; it is systematic; and it is available to us by the use of reason.

Optimists' View: Ethics is (somewhat) like mathematics

But: (i) mathematics does have a pretty impressive track record; think of the sheer body of mathematic results, on which there is pretty much complete agreement (there is some dispute over, e.g. what set theory to use, but that seems more like disagreement in application than in theory). Work on ethics started around the same time, but it doesn't come close. (ii) mathematics may not be empirical, but it is closely tied into the empirical sciences; in fact it is indispensable for them. Ethics doesn't seem to be tied into the empirical sciences in the same way. It looks conceptually isolated: the fact/value (is/ought) distinction. Is it indispensable?

One reaction to this is to give up on ethics:

Pessimists' View: Ethics is illusory

Two forms of this: one actually rejects our normative ethical views. We are all really egoists, and ethics is just rhetoric (or perhaps self-deception, or a false ideology). Alternatively: revise our meta-ethical error theories: we are radically wrong about the *nature* of ethics; it has no objective basis. Doesn't the second entail the first? Quite a lot of error theorists want to hang on to the normative ethical views, while being metaethical error theorists (e.g. Mackie), but it's unclear that that is really available. And can we really contemplate giving up on ethics altogether? Isn't it *practically* indispensable.

There are also metaethical views that aren't so pessimistic, but still hold that our pre-theoretical views are radically wrong. Emotivism, and its generalized version, Expressivism, fall into this class. They deny that moral judgments are actually *judgments* at all, while saying that it is fine for us to go on using moral sentences. But they face serious difficulties too. And again it is far from clear that they are consistent with maintaining our current ethical practices.

D. A third possibility

Consider linguistics. Linguists also make use of intuitions; their theories are meant to account for them. And the theories linguists come up with are typically complex (more complex than most existing normative ethical theories). Psychology might help with linguistics; psycholinguistics is increasingly important in understanding how our language works. (Note the difference here with mathematics. There is an interesting psychology of mathematics: findings on how street traders manage to do complex calculations in their heads for instance. But we wouldn't expect to find anything about the mathematics itself from psychology.) So that suggests a third possibility:

Moderate Optimists' View: Ethics isn't illusory, but it's more like linguistics than like maths

Linguistic theories are typically complex and messy, full of exceptions (that is why it has been so hard for AI to replicate human speech; and why machine learning has done better than programmes that try to explicitly state the rules). And linguistics is concerned with linguistic behaviour generally, not just with judgments: sometimes our judgments about what we say come apart from our actual linguistic behaviour. Consider this sentence (with this punctuation):

The horse raced past the barn fell.

Is it grammatical? This one clearly is:

The mouse the cat killed ate the cheese.

But it is less clear that this one is:

The mouse the cat the dog chased killed ate the cheese.

If we can be wrong in our intuitive judgments about which sentences are grammatical, then we might think that shows there is a difference between what we *do* say and what we *ought* to say. In that sense then linguists might be thought to have some normative power. If we liken ethics to linguistics, is that enough to explain its normativity: to explain why ethics has a genuine authority over us? It's not just that ethics places demands on us (you might think that grammar places them to); it is that we feel them to be in some sense *categorical*, though saying quite what this amounts to is no easy job. At least this seems right: it would be very odd to convince someone that something is morally required, and for them to not be in the least motivated to do it (or at least, not to feel guilty if they didn't, unless they display genuine psychopathic tendencies). Ethics here looks to display what is often called 'motivational internalism' (though beware that 'internalism' is used as the name for many different doctrines in this area). It would be odd to think that we'd have made a huge mistake if we spoke a language in which 'cat' meant dog, and 'dog' meant cat. But we are not similarly relaxed about ethics. We don't think it would be fine if we embraced a theory in which, say, racism was admirable. So is there any way of seeing linguistics as providing a model in some ways, but not in others?

Clearly these themes combine questions about metaethics with questions about moral psychology: what sorts of mental states our moral judgments really are, whether they motivate us and if so how, and so on. Our aim will be to try to make some headway by investigating both. We'll start with one pessimist view: the view that ethics must be deeply illusory, since we are all bound to be egoists.