

Week Thirteen: Moral Realism

1. Ramsey Sentences and Folk Morality

We looked last week at one toy example of terms being defined by just using them rather than by, say, giving a stipulative definition, or a baptism, or any of the other standard ways to define terms. This was the detective example, where the detective simply gives names to the fillers of some functional roles in the crimes. If the theory of the crime is roughly true then those who fill those roles really get those names.

As we saw, the Ramsey-Lewis-Jackson method for defining terms whose philosophical status is obscure is parasitic on this approach. We presume a starting point where we don't understand the words of some disputed area, say the mental or the moral, but we know how some people use those terms. In the cases in which we are interested we don't look for a single detective, but take the uses of a large body of people. Since these people differ (radically in the case of morality) about their theories, this adds a layer of complication, and maybe that layer will become important soon. When all goes well, we can use the same method as we used in the detective case; if the theory about the mental or the moral (folk psychology or folk morality) is roughly true, then the disputed terms will name some parts of the world which serve as truthmakers for that story.

Several complications immediately ensue. First, how much should we worry about the fact that the terms in question weren't introduced by this method, the myth of our Rylean ancestors, really is a myth? There is little in the literature on this question, and I'm not sure what to make of it. Secondly, is the myth even coherent, or is it impossible to grasp the ur-concepts without grasping some of the controversial concepts as well? (I think Brandom thinks this is impossible, but it really is ill-advised to rely on my accounts here.) And if it is impossible, does that make the theory of meaning any less plausible. We talked a little about these last week, so I'll spend more time in these notes on other problems.

Third problem. The folk don't have a unified theory of morality, so what should we take when we ask what folk morality really is. If we took an opinion poll, most folk's theory of morality would closely follow that found in one or other religious texts, such as the Bible or the Koran. But right-thinking folk (meaning, our friends) agree that these theories are fundamentally wrong-headed. So the right analogy is with the case where the detective introduces new names by a theory of the crime which is wrong in all important respects. The consensus seemed to be that in such a case the new names don't refer, so we are led to an error theory of ethics.

Maybe this is a reason to have the defining story be mature folk morality, rather than current folk morality. Remember that in Jackson's account we don't Ramsify over the current ethical platitudes. Rather we Ramsify over an idealised version of folk morality. And there seems to be some consensus that mature folk morality will be atheistic. (This isn't a consensus among the folk, of course!)

This isn't the reason that Jackson gives for Ramsifying over mature folk morality, rather than current folk morality. Rather the reason he gives is that current folk morality is full of unresolved tensions, and we should fix those before we work out the reference of moral terms. I really don't understand the argument here. We noted already that a term-introducing story can be 'successful', the terms it introduces can refer, even though the story is false. Why can't we say the same thing about current folk morality? Granted it is probably false. After all it is probably inconsistent. But there may be enough truth in it for it to successfully introduce terms.

I'm not sure that it is *possible* to Ramsify mature folk morality rather than current folk morality, but I suspect I'm just misguided here. Just what is it we do when we try and make folk morality grow up? We try and find which parts of it are more central than others, by seeing which parts of it we are most attached to when faced with the conflicts between its parts. (This is a premise; I might have the phenomenology of ethical research all wrong here. But this is what it feels like to me.) But this is just what we do when Ramsifying current folk morality. Recall that when there are no perfect deservers, the term-introducing story isn't true in all respects, we have to work out which parts of the story can be sacrificed before the theory fails to introduce new terms. So it seems that all we can do is what we would do were we Ramsifying current folk morality.

I suspect one reason people have for disliking this emphasis on current folk morality is that there is no guarantee that Ramsifying it will yield a unique reference. So consider the possibility that each of the reference assignments 'good acts' = 'acts in accordance with duty' and 'good acts' = 'utility maximising acts' preserve enough of folk morality for it to be a successful term-introducing theory. (As Adam said last week, the easier we make it for terms to refer at all, the more likely we make it that they will not *uniquely* refer.) If this is so, and I suspect it is, there is no fact of the matter whether Kantian or utilitarian ethical theories are true. The moral terms which are used to state those theories are simply ambiguous between deontological and utilitarian concepts. Well, that's not a disaster to me (it looks *prima facie* true), but I don't have much emotional investment in either side of this debate!

Last week there were some good questions about how Jackson's approach makes it contingent that there are moral properties. In our toy example about the robbery, the term-introducing theory was not guaranteed to be even approximately true, and if it wasn't even approximately true then the terms did not refer. By analogy, it might be expected that there is some

risk that when we plough through the moral platitudes to find the Ramsey sentence for ethical terms, that sentence will turn out to be so badly wrong that the terms it purports to introduce do not refer.

The analogy here is imperfect because we are introducing terms which name properties not individuals, and the existence criteria for properties are much less clear than they are for individuals. So take a term-introducing story which does seem to be badly wrong – the phlogiston theory. On that story, phlogiston is the stuff which is given off during burning. It turns out that nothing is given off during burning, burning is a matter of taking in oxygen, not giving off anything, so there is nothing which is phlogiston. But it isn't clear that the term 'phlogiston' doesn't denote, it seems plausible to say that the term denotes a property which is not instantiated. The same may well be true of the moral terms; like 'phlogiston' they are sure to denote some property, but it is contingent whether that property is instantiated. And of course this is true; in rock worlds no moral properties are instantiated.

Just as there is a difficulty in giving existence criteria for properties, there is some difficulty in giving identity criteria. So it isn't clear to me how Jackson can distinguish (as he purports to do on pg 141) between role properties and realizer properties. More interestingly, we might be able to dissolve some disputes in ethics if we are liberal in interpreting claims of property identity. So as Jackson notes, the Ramsey sentence approach opens the pleasant possibility that virtue ethicists and utilitarians are both right in some sense. It may be that rightness is the property distinctive of acts of the virtuous, but in practice the property distinctive of acts of the virtuous is utility maximisation. Again, this only seems plausible if we like to dissolve ethical debates rather than continuing them, but that isn't a problem for me!

2. Permutation Problem

Michael Smith argues that the Ramsey sentence approach cannot work because it is vulnerable to what has been dubbed (by Philip Pettit) the 'permutation problem'. The way Smith argues for this is to argue that a similar problem besets any Ramsey-style definition of the colours, and then arguing that the same argument applies to definitions of moral terms.

Let's say we were to use the Ramsey sentences to define our colour terms. First we have to collect the platitudes about the various colours. So the platitudes will include things like 'Red is more similar to orange than it is to blue', 'Red objects look red under standard conditions' and so on. To create a Ramsey sentence, we have to factor out all the colour words from each of the platitudes. So these platitudes become 'There are colours x_1 , x_2 and x_3 such that x_1 is more similar to x_2 than it is to x_3 ', 'There is a colour x_1 such that objects which have colour x_1 appear to have the colour x_1 ' and so on. Smith's claim is that none of these platitudes will serve to individuate the colours. For every

platitude about 'red', there is a symmetric platitude about 'blue'. This follows from the fact that the colours are all on a symmetric colour wheel. So we need something more than platitudes to get the reference of the colour words.

Return to our detective story, about Marlowe giving his theory of the crime. Imagine that if he never said that Lefty did one thing and Help another, so every sentence was like "Then Lefty and Help went into the bedroom and attacked Miss Smith." Even if the story is true, and there are two robbers who do just what Marlowe said Lefty and Help did, there would be no fact of the matter which robber was Lefty and which Help. There is nothing which could make this robber Lefty and this robber Help. This is what Smith claims happens in the case of the colours.

Of course there are other platitudes which could be used to fix the reference of the colour words. If we add platitudes like 'Red is the colour of blood', 'Yellow is (normally) the colour of a new-born chicken', the story will say enough to individuate the colours. But Smith claims that only *a priori* platitudes can be used in the reference-fixing story. The argument for this is not entirely clear. I think the reasoning is as follows. What is needed to grasp a concept is knowledge of the *a priori* relations between it and other concepts. The term-defining story should only include platitudes needed to grasp the concepts underlying the defined terms. I don't really see why either premise should be accepted, but we'll leave that to one side. There is another argument which could be given. (And maybe I'm being unfair to Smith and this is the argument he does give.) Facts about meaning are knowable *a priori*. The meaning-determining platitudes are, in this sense, facts about meaning. So only things knowable *a priori* can be platitudes. So maybe this step of the argument can work.

A larger problem with the argument at this stage is that there are platitudes, even *a priori* looking ones, which possibly can be used to differentiate the colours. It just isn't true that the colour wheel has the kind of symmetries that Smith needs for the argument. To take the most obvious example, there are no dark yellows, though there are dark reds and dark blues. Maybe these are *a priori* in some sense; can we imagine a possible world in which they are false? These distortions in the colour sphere, if admissible as platitudinous evidence, are sufficient to uniquely pick out the colours. (This is argued in a recent paper in the *Australasian Journal of Philosophy* by Alan Hazen.)

If we accept that only *a priori* platitudes are allowed, and join Smith in what we take to be *a priori*, we can possibly generate a permutation problem for longitudes. I'm not entirely sure how to analyse the sentence, "Syracuse is at 43°N, 76°W", because I don't know what the semantics for latitudinal and longitudinal measurements should be. (Are they names, descriptions, what?) In theory, I should be able to use the Ramsey sentences to help me. But when we look at all the *a priori* platitudes about latitudes and longitudes, all we find are relative platitudes: that 76°W is a little west of 75°W (and a lot further east!), that 43°N is north of 42°N, that latitudinal lines are of different

lengths, but longitudinal lines are of (roughly) the same length, and so on. The platitudes may fix that the equator is at 0°N . But none of the platitudes will fix which line is 0°W , so knowledge of the platitudes about longitudes will not tell us where in the world 0°W is. This seems a better candidate than the colours to serve as an exemplar of the permutation problem. Even though there are deformities in the earth sphere as well as in the colour sphere, there is no *a priori* connection between the deformities in the earth sphere and the longitudes, but there *maybe* an *a priori* connection between the deformities in the colour sphere and the colours.

When a permutation problem arises, the difficulty for the Ramsey sentence method isn't just that the platitudes are multiply realised. As we saw last week, multiple realisation is perfectly compatible with the terms referring, as long as we are prepared to live with ambiguity. (And why shouldn't we be?) The real difficulties are that (a) there are too many different resolutions; and (b) some of the different resolutions are polar opposites. These do seem to be real concerns. But it isn't clear why we should think this problem arises for the moral terms. Smith gives three reasons for thinking so, but they all seem to be pretty bad reasons. And there are reasons for thinking that the *a priori* platitudes are sufficient to fix the reference of the ethical terms.

Smith's first reason¹ for thinking a permutation problem will arise is that we learn the ethical terms by being presented with paradigms, rather than by learning the *a priori* platitudes about morality. But we also learn the concepts *triangle*, *cube* and so on by being presented with paradigms, rather than learning the platitudes. (As the history of solid geometry shows, it is non-trivial to get the *a priori* facts about polyhedrons exactly right, so this method of acquaintance is probably pedagogically wise.) And no one thinks there is a permutation problem for the geometric concepts.

The second reason for believing that there is a permutation problem is that it would explain the abject failure of attempts to provide a naturalistic definition of the moral concepts. But that failure can also be explained by the existence of a non-vicious ambiguity in the moral concepts, such as an ambiguity between the concepts promoted by deontologists and utilitarians. (Many would argue about the viciousness of such an ambiguity, but it seems fine to me.) So this doesn't seem like a reason to postulate a problematic permutation problem; it isn't as if under one resolution it turns out that starving the peasants is morally acceptable.

The third reason Smith provides is that the *a priori* platitudes are interconnected. If every platitude has several moral concepts in it, this would provide some reason for worrying about a permutation problem. But I doubt this really is the case. It does seem *a priori* that causing pain is bad, other things being equal. As far as I can see, this only has one moral concept in it, unless it turns out

¹ The ordering here is kind of arbitrary.

that the *ceteris paribus* clause is too morally loaded. It does seem to me that there will be enough to ground the moral concepts, but maybe I'm just being naïve.

Perhaps a better example of this kind of grounding can be found in the epistemic platitudes. Remember that we are trying to define all the normative concepts at once by means of a long Ramsey sentence, so presumably we are trying to define the epistemic norms as well. And there do seem to be enough *a priori* platitudes about the epistemic to avoid any permutation problem. For example, it seems platitudinous that using *modus ponens* in inferences is rationally permissible. (I know there are those who claim *modus ponens* is really invalid; if you agree substitute your preferred obviously valid inference.) So there is no permutation problem for the epistemic norms. If there are enough interconnecting platitudes to fix the moral norms given the epistemic norms (and I see no reason to suspect otherwise) there is no permutation problem for the moral concepts.

So in sum, Smith's argument relies on two steps, each of them dubious. First, he argues that we are only allowed to use *a priori* platitudes to fix the reference of the moral concepts. Secondly, he argues that the *a priori* platitudes do not provide enough constraints on the moral concepts to provide anything like a determinate reference. So the Ramsey sentence approach is wrong. Neither premise seems compelling, and the second seems more likely false than true, so the argument is not in good shape.

3. Cornell Realism

Jackson sets up his version realism by contrasting it with what he calls 'Cornell Realism', so-called because of its spiritual home in central NY. Cornell realists agree with Australian realists that ethical properties are identical with descriptive properties, but we cannot analyse ethical sentences in descriptive terms. An extended statement of this position is found in Boyd's "How to be a Moral Realist", which was on the readings. And a quick statement is from Nicholas Sturgeon's "Moral Explanations", from which I've attached an extract.²

There appear to be two points of disagreement between the Cornell crew and the Australian crew. I'll try and present each of these neutrally, but I should point out that on each I think there is a clear winner. The two issues are (a) whether there are limits to the lengths of analyses, and (b) whether we need to do *a posteriori* investigation to determine the nature of the *concept* of goodness.

² Both papers are in Sayre-McCord (ed) *Essays on Moral Realism*.

3.1. *Limits to analyses*

In Jackson's 'proof' that there must be something like an analysis for the ethical terms, he assumes that there are no limits to the lengths of sentences. So it is perfectly acceptable if 'good' is equivalent to a long disjunction of descriptive clauses, where there are as many disjuncts as there are possible worlds. As Sturgeon points out, this is impossible in anything which resembles languages we work with. It is normally assumed that there are at most countably many terms in a language, and that sentences have finitely many clauses (or maybe countably many). But there are several more possible worlds than this.

Digression on infinities. It is standard to say that two sets are the same size if there is a one-one correspondence between their members. The various infinities are identified as the sizes of some well-known sets. The plural 'infinities' here is appropriate because there are infinite sets which are not the same size according to this conception. Cantor proved that there is no one-one mapping between the natural numbers and the reals. (The neat part of the proof is a method for generating a real which is left out of any proposed mapping.) The proof doesn't just rely on the fact that there are reals which are not natural numbers, but not conversely. There is, for example, a one-one mapping between the naturals and the even naturals: 1 to 2; 2 to 4; 3 to 6; and so on. Anyway, a countable set is one which can be mapped onto the naturals. *End of digression.*

Three possible outcomes. First, Sturgeon is right that there are limits on the size of what is properly called an analysis. Secondly, Jackson is right that we can call something a reductive analysis even if it 'reduces' a sentence like 'Baseball is evil' to a disjunction of uncountably many disjuncts. Thirdly, there is just no fact of the matter here, and we are just free to stipulate what counts as an analysis and what doesn't. For my money, the second looks like a non-starter, the third looks possible, and the first looks probable.

This point about the infinitary nature of the full story about the world might help explain some of the oddity of some of what Jackson writes. So on page 147 he says, "Metaphysical descriptivists think that how the world is, how we take things to be, and conventions of word usage, can be exhaustively given in purely descriptive terms." Well, this isn't clearly true. Metaphysical descriptivists think that there are no facts over and above the descriptive facts, but if we think there are limits to the size of story we can tell, or to the number of primitive terms in that story, this alone won't entail that we can do what Jackson says is possible. This oddity seems to me to be a point in favour of Sturgeon's account. But maybe the right thing to do here is to say the question is one which can be settled by stipulation.

3.2. *Morality and Contingent Facts*

The other half of this debate is much harder to assess. Here's what looks to be Boyd's argument, though I must confess I may have this totally wrong. The only way we could give a reductive analysis of moral terms is if the relationship between the descriptive and the ethical were *a priori*. But it is *a posteriori* which concepts are the moral concepts. For example, it may be that it is good to promote friendship. This may turn out to be contingent: in some worlds people may be better off without friendship.

Boyd tries to back up this argument by a comparison with the futility of trying to analyse causation. We should try and find out what causation is, he claims, by doing some empirical investigation. This just must be a mistake, at least on one level. Imagine a fantasy story where one character casts a spell, and another character immediately turns into a frog. It seems plausible to say that the spell caused the frogginess. But there is no property instantiated in this world which is instantiated in that part of the story (except 'frogginess'). So causation can't be some property we can discover by empirical investigation.

The same is presumably true for goodness. We can tell fantasy stories where there are clearly good and evil characters, even though none of the properties Boyd provisionally identifies on page 203 are affected by actions of either side. So goodness can't be identified with the promotion of those properties, even *a posteriori*.

There is a sense in which goodness and maximising Boyd's style of flourishing may be identical. It may be that in practice every good act is an act which maximises flourishing. But this just shows that the two properties are co-extensive. And if we say that is an identity, we seem forced to say that the properties 'being a creature with a heart' and 'being a creature with a kidney' are identical, which seems odd. So I don't get where this argument is headed. Even once we've identified some descriptive properties whose promotion is unequivocally good, we need to be able to (a) convince the sceptics that promoting these properties really is good, and (b) say what would be good in various counterfactual situations. I don't see how observation, unaided by any kind of analysis, will be helpful here.

Even if we ignored all that, there is a further problem for this style of argument. As Jackson points out, even if it is *a posteriori* which property plays the goodness role in this world, this is compatible with the existence of his kind of analysis. What would really be problematic is if the inference from the full descriptive story to the ethical story was not *a priori*. And Jackson argues (briefly) that there is no reason to think this. The point is just the one we made about the descriptive story entailing the water story, because it entails (*a priori*) what thing fills the water role.