A problem for expressivism

Frank Jackson & Philip Pettit

1. Introduction

Language, Truth and Logic added expressivism to the inventory of substantive positions in meta-ethics, and the recent defences of versions of it by Simon Blackburn and Allan Gibbard have enhanced its status as a major position.[1] Ayer presented the doctrine as an improvement on subjectivism – that is, on the doctrine that ethical sentences serve to report attitudes of approval and disapproval – and it is widely supposed to be an internally coherent and interesting position. We argue, however, that there is a serious problem that expressivists, unlike subjectivists, have to face which has not been adequately addressed in the development of the doctrine.

Expressivism is a bipartite theory. It holds, first, that ethical sentences lack truth conditions – they are not truth apt, truth assessable etc. – and do not serve to report anything that the speaker believes to be so. And it holds, second, that ethical sentences express certain distinctive pro and con *attitudes*, and express them without in any sense reporting them: without reporting them even in the broad sense in which 'The present King of France is bald' reports the existence of the King as well as his baldness.[2] The first clause is a negative, semantic claim; it denies that ethical sentences have truth conditions. The second is a positive claim: it informs us about the function – that of expressing attitudes – served by ethical sentences.

The second, positive clause is supposed to underpin the first, semantic claim in the following sense. If ethical sentences reported the presence of

[1] A. J. Ayer, *Language, Truth and Logic*, London: Gollancz, 2nd. ed. 1946, Simon Blackburn, *Spreading the Word*, Oxford: Clarendon Press, 1984, and Allan Gibbard, *Wise Choices, Apt Feelings*, Oxford: Clarendon Press, 1990.

[2] For a view according to which ethical sentences report attitudes but only in this wide sense – in a sense in which the presence of the attitudes is not the focus of attention – see James Dreier, 'Internalism and Speaker Relativism', *Ethics* 101: 6–26.

Analysis 58.4, October 1998, pp. 239–251.

attitudes rather than merely expressing them – if they expressed beliefs that the attitudes were present – then, problems of vagueness aside, they would automatically be true or false. They would be true just when the beliefs they expressed were true, just when the attitudes they reported obtained. The reason, according to expressivists, that ethical sentences lack truth conditions is precisely that they express attitudes without reporting them.

The problem that we raise for expressivism is that expressivists do not have a persuasive story to tell about how ethical sentences can express attitudes without reporting them and, in particular, without being true or false. We think that they cannot satisfactorily explain how the two clauses in their theory can obtain together: how, without being true or false, as the first clause maintains, ethical sentences can yet serve to express certain attitudes, as the second clause insists that they do. Although there are (of course) interpretations that might be given to the claim that ethical sentences express attitudes which would mean that ethical sentences lack truth conditions, we shall see that they are not interpretations that expressivists can afford to embrace.

Most of our paper is given to presenting this problem but a final section is devoted to the issue, as it is often described, of the persistence of moral disagreement. Expressivists argue that, precisely because they deny truth conditions to ethical sentences, they are especially well-placed to make sense of persisting moral disagreement. This, indeed, is the argument with which Gibbard launches *Wise Choices, Apt Feelings*.[3] In the final section, we argue that this is a mistake: the denial of truth conditions to ethical sentences is not merely hard to maintain for the reasons we will be presenting, it does not buy a solution to the puzzle that so often motivates it.

One further remark before the main business. We said above that expressivism holds that ethical sentences serve to express attitudes. This is a simple version of expressivism. More complex versions might talk of expressing higher-level attitudes, or attitudes formed after due deliberation or negotiation with others, or the properties to which the attitudes are expressed, and so on. In the same way, we can distinguish simple versions of subjectivism that hold that ethical sentence report the attitudes of the utterer from more sophisticated (and plausible) versions that focus on higher level attitudes, properties of actions inasmuch as they prompt attitudes in us and others, and so on. We shall continue to frame our

[3] See chapter 1, esp. pp. 11–18; see also Blackburn, *Spreading the Word*, p. 168, and chapter 1 of C. L. Stevenson, *Ethics and Language*, New Haven: Yale University Press, 1944.

discussion in terms of the simple versions. The argument is more easily developed in relation to them and extends *mutatis mutandis* to the more complex ones.

2. *Locke on voluntary signs*

The crux of the problem we raise for expressivism is that a point of Locke's about language makes it very hard for expressivists to hold onto their view that ethical sentences express attitudes in a way that prevents those sentences from having truth conditions.

In holding that ethical sentences serve to express attitudes, expressivists do not suggest that this occurs in an unintentional or accidental manner: in the manner in which a natural response like a gasp may express admiration, or a yawn boredom. They regard words and phrases like 'good', 'ought to be done', 'just', and 'morally bad' as conventional signs which we learnt when we learnt English. And they maintain that when we use such terms we intentionally express the attitudes of approval and disapproval with which they are associated.

Now there is a serious problem with the claim that certain words and expressions are conventional, intentionally used signs and yet generate sentences that lack truth conditions. The problem goes back to Locke's observations about voluntary signs.[4]

Locke observes that, because it is contingent and fundamentally arbitrary that we use the words we do for the things we do use them for, our ending up with the conventions or arrangements we have in fact ended up with is to be understood in terms of our, explicitly or implicitly, entering into agreements for the use of these words for these things.[5] However, entering such agreements requires that we *know* what it is that we are using the words for. As Locke puts it, 'Words being voluntary signs, they cannot be voluntary signs imposed by him *on things he knows not*. … they would be signs of he knows not what, which is in truth to be *the signs of nothing*.'[6] In other words, because the word 'square' is a voluntary, agreed-on conventional sign of the property we use it for, the explanation of how it is that we come to use it for that property requires that, on certain occasions, we *take* something to be square, and use this fact to found the convention of using the word 'square' for that property. To make the point

[4] John Locke, *An Essay Concerning Human Understanding*, book III, chapter 2.

[5] For more on this, see, e.g., see H. P. Grice, 'Meaning', *Philosophical Review*, 66 (1957): 377–88, David Lewis, *Convention*, Cambridge, Mass.: Harvard University Press, 1969, and Jonathan Bennett, *Linguistic Behaviour*, Cambridge: Cambridge University Press, 1976.

[6] *Essay*, III.ii.2. Our emphasis.

vivid, consider your reaction to being asked to enter an agreement to use the word 'Fred' for a certain shape by someone *who won't let you see which shape it is*.[7]

If this line of thought is right, then any explanation of how we English speakers came to use the voluntary sign 'good' for the attitude we do use it to express, according to expressivists, must allow that we *recognize* the attitude in question in us. For, to follow Locke, we could hardly have agreed to use the word for an attitude we did not recognize and failed to believe we had, since that would be to use the word for 'we know not what'. But that is to say that expressivists must allow that we use the word sincerely only when we *believe* that we have a certain kind of attitude. And then it is hard to see how they can avoid conceding truth conditions to 'That is good', namely, those of that belief. Not only will the sentence 'That is good' express the attitude alleged, it will be true just in case the attitude is present and false otherwise: it will in some sense, however broad, report the presence of that attitude. Expressivism will have become a variety of subjectivism.

Here is a way to put the point. The attitudes that expressivists hold are expressed by the use of ethical terms in sentences had better be attitudes we (and our potential audiences) recognize as those it is correct, according to the conventions of English (or French or …), to use the ethical words for, and it had better be that we use the ethical words just when we believe that we have the attitudes. We are not ostriches when we give voice to our attitudes. Moreover, when we follow the conventions for the use of the ethical vocabulary, we seek to conform to them, and know, or at least very much hope, that those we are communicating with are seeking to conform to them. We regulate our linguistic behaviour under the presumption that our fellows are doing likewise as part of our membership of a common linguistic community. Would it be better, according to expressivists, to use the ethical vocabulary when we thought we had quite different attitudes, or had no opinion at all on the matter, and would it be better if we ignored the usage patterns of our fellow language speakers? How then, runs our challenge, can *expressivists* fail to count our uses of the ethical terms in sentences like 'That is right' as reports of those attitudes?[8]

The view that ethical terms serve to report attitudes is compatible, it should be noted, with the view that they express the attitudes. Indeed, a

[7] The issue of whether we start with words or (one word) sentences does not affect the points we make here; the key point is simply that we perforce start out with *belief*. According to Lewis, *Convention*, in order to found the convention, we need much more than the belief that something is square; we need a whole web of beliefs about beliefs including those of our fellow speakers of the language. For our purposes here, the crucial point is that we need *at least* to regulate our use of the word via belief.

good way to express an attitude is to report that you have it. It would be absurd to tell someone who says 'I dislike George a good deal' that they failed to express their dislike because they reported it. But if 'That is right' and the like are reports as well as expressions of the attitudes, then they have truth conditions, contrary to the central claim of expressivists.

Ayer half-sees the difficulty that this observation raises for expressivists in his discussion of a 'complication' concerning the crucial distinction between expressing and asserting.

> The distinction between the expression of feeling [his term for the relevant attitudes] and the assertion of feeling is complicated by the fact that the assertion that one has a certain feeling often accompanies the expression of that feeling, and is then, indeed, a factor in the expression of that feeling. Thus I may simultaneously express boredom and say that I am bored, and in that case my utterance of the words, 'I am bored', is one of the circumstances which make it true to say that I am expressing or evincing boredom. But I can express boredom without actually saying that I am bored. I can express it by my tone and gestures, …, or by an ejaculation, or without uttering any words at all.[9]

Ayer sees that saying that one has a certain attitude counts as expressing it but observes (rightly) that there are ways of expressing an attitude that do not involve saying that you have it. But his examples are ejaculations and gestures, and cases that do not involve 'uttering any words at all'. This hardly helps with sentences like 'This is good' and 'That is wrong', which involve words, are not typically ejaculations, and, moreover, are part of the voluntary convention we English speakers entered into concerning what words stood for what when we mastered the language.

Blackburn also notes that one way of expressing an attitude is to say that you have it.[10] He is explicit that his view is not that expressing attitudes,

[8] The issue here is quite separate from whether the attitudes reported are themselves in the general category of beliefs. Expressivists are typically internalists and Humeans – they argue that moral judgements have internal connections to motivation that Hume taught us no cognitive state can have. This would mean that moral judgements are not beliefs proper but that does not negate our point that ethical sentences had better, in that case, express subjects' beliefs about their allegedly non-cognitive judgements. This point was overlooked by one author (FJ) in Frank Jackson, Michael Smith and Graham Oppy, 'Minimalism and Truth Aptness', *Mind*, 103 (1994): 287–302.'Minimalism and Truth Aptness', but he stands by the critique of minimalism about truth aptness in that paper.

[9] *Language, Truth and Logic*, p.109.

[10] Blackburn, *Spreading the Word*, pp. 169–70. In his view, you can, in addition, express attitudes by making claims about how things are other than the claim that you have the attitude in question.

on the one hand, and reporting them, describing them or saying that you have them, on the other, are incompatible; it is simply that you can have the former without the latter. We agree. The problem is, as we have been arguing, that the cases where you can have the former without the latter differ in a crucial respect from the way we use ethical words and sentences.

It might be objected that the production of a sentence like 'That is wrong' is much more like an exclamation – or ejaculation, as Ayer puts it – than we are allowing: we contemplate some situation and simply find the words 'That is wrong', as it might be, coming to our lips. The attitude 'outs' itself without any conscious cognitive processing, and so, expressivists might argue, there is no truth-assessable claim being made when we produce the sentence. What happens is almost like wincing. This seems to us a highly implausible account of our moral psychology. It is a commonplace that we argue about and reflect on which ethical judgements to make. We do not typically 'give out' with the first thing that pops into our head. This is a point that Blackburn is explicit about when he talks about 'taking *all possible opportunities* for improvement of attitude'.[11] Moreover, even if some relatively simple kind of 'outing' story were true, this would not support the view that ethical sentences lack truth conditions. Face recognition is mostly not a conscious cognitive process, which is why verbal descriptions of suspects are often of little use to the police. Nevertheless, judgements of facial identity have truth conditions – obviously. In the same way, even if ethical sentences are prompted by an ability to recognize certain attitudes, they would still have truth conditions; they would simply be claims about the recognized attitudes.

This completes the presentation of our main challenge to expressivists. In the next three sections we consider three replies that they may (do) make and try to show that none is satisfactory.

3. First reply: expressing versus reporting belief

Expressivists, and philosophers in general, often rightly distinguish expressing what you believe from reporting what you believe. The sentence 'Snow is white', uttered in the right context, expresses your belief that snow is white, and is true iff snow is white. It does not report your belief. If you want to report your belief that snow is white, you need to use the sentence 'I believe that snow is white' (if you are speaking English), and this sentence is true iff you have the belief.

Expressivists often suggest that we can apply this distinction to other psychological states, including especially the 'ethical' attitudes, and that when we do, we get the account they need of the sense in which ethical

[11] Blackburn, *Spreading the Word*, p. 198, our emphasis.

sentences express attitudes. They observe that we can distinguish the doctrine that 'X is right' reports a certain pro-attitude to X, from the doctrine that it expresses that pro-attitude to X. The first view is subjectivism; the second, they claim, is expressivism.[12] They may argue, then, that the availability of this distinction shows that there has to be something wrong with the argument from Locke. There has to be a sense of 'express' which ensures that 'X is right', and the like, express attitudes without reporting them and without having truth conditions.

Sometimes it almost seems as if the argument on offer here is supposed to run as follows: if 'X is right' reports an attitude, 'X is right' has truth conditions; therefore, if it expresses an attitude instead, it lacks truth conditions. But this would, of course, be to commit a fallacy: a version of denying the antecedent. What we need to do, in order to give the argument a fair hearing, is ask about the sense in which 'Snow is white' expresses a belief, apply this sense to the claim that 'X is right' might express an attitude, and see whether that will really help expressivists. We maintain that it will not.

The trouble for expressivists is that, although there is an important difference between reporting and expressing a belief, it is plausibly a difference in what is reported. It is not a difference between reporting something and not reporting at all. When you express your belief that snow is white by producing, in the right context, the sentence 'Snow is white', you are not avoiding the business of reporting altogether. You are not reporting the fact that you believe that snow is white but you are reporting the content of that belief; your are reporting that snow is white. This is how the sentence gets to be true iff the belief is true.

If we take the distinction as drawn for beliefs, then, and apply it to attitudes, we do not get a result that can help expressivists. What we get is that 'X is right' expresses a certain pro-attitude iff 'X is right' reports the content of the attitude. And this is not at all what expressivists are after. First, it makes 'X is right' out to have truth conditions, namely, those of the content; and, second, it is very implausible in itself. The relevant content will something like *that X happens* – for that is what we are favourably disposed towards, according to expressivists, when we assert that X is right – and that is very different from X being right (unfortunately).

4. Second reply: the exclamation analogue

There are plausible examples of intentionally delivered, more or less conventional signs of psychological states that seem to lack truth condi-

[12] See, e.g., Blackburn, *Spreading the Word*, p. 169, and Gibbard, *Wise Choices, Apt Feelings*, p. 84.

tions and both the second and third replies to the argument from Locke take their starting point from these. The second reply starts from examples that expressivists themselves often cite: exclamations like 'Boo' and 'Hurrah'. Such exclamations do not have truth conditions and are not used to make claims about how things are – so at least it appears – and yet they are conventional, intentionally employed signs of psychological states. The suggestion often made by expressivists is that 'X is good' is very roughly like 'Hurrah', and 'X is bad' very roughly like 'Boo'. In particular, these sentences are like the relevant exclamations in not having truth conditions and so in not making claims about how things are.

We are prepared to grant that exclamations like 'Boo' and 'Hurrah' do not have truth conditions. But we insist that that possibility very much calls for explanation, particularly in the light of the argument from Locke. And the only available explanation, we will argue, is not one that can be plausibly extended to explain how ethical sentences might lack truth conditions.

Even if 'Boo' and 'Hoorah' do not have truth conditions, they could have had them. We could, for example, have used 'Boo' simply to mean that the producer of 'Boo' does not like what they are here and now acquainted with, in which case 'Boo' would have had truth conditions; it would have been true in S's mouth at *t* iff S did not like what S was acquainted with at *t*. What then could stop 'Boo' from having these, or indeed any, truth conditions? The reason does not lie in the fact that 'Boo' lacks the right syntax – it would be strange if syntax alone stopped us making claims about how things are. The only possible explanation is that, as it happens, 'Boo' has a rather erratic, highly personal connection to our feelings.

The situation we have in mind is like that which obtains with the word 'but'. It is standard doctrine that 'P and Q' and 'P but Q' have the same truth conditions, but it might, of course, have been the case that their truth conditions differed. It might have been that the contrast signalled by 'but' was part of the literal meaning of 'P but Q', in which case the existence of a contrast would have been a necessary condition for the truth of 'P but Q'. The reason it isn't is that the convention of using 'but' when there is a contrast is not sufficiently entrenched and clear cut.[13] (And, as you would expect if this is right, the minority who think that 'but' does affect truth conditions are those who insist that the convention *is* sufficiently well-entrenched and clear cut.)

Why does the convention need to be so well-entrenched and clear cut? This is a controversial issue in the philosophy of language, but the essential point is reasonably uncontroversial. When we produce sentences to say

[13] This account of why 'but' signals a contrast without altering truth conditions is essentially the one Michael Dummett, *Frege*, London: Duckworth, 1973, offers on p. 86. Incidentally, as he points out, the contrast signalled need not be between 'P' and 'Q'.

how we take things to be, there will typically be a whole web of beliefs lying behind their production. The problem is to isolate the ones that matter for what is literally said, in the sense of determining truth conditions, by using the sentences. For instance, typically, an English speaker only says 'Snow is white' when they believe that there is someone around to hear what they are saying, but that is not part of what they intend to convey and is not a necessary condition for what they say being true. It is the nature of the complex set of conventions for the use of the sentence that serves to filter out the 'wrong' candidates. We are using the phrase 'well-entrenched and clear cut' to mark out the conventions that do the trick. For example, followers of the 'Gricean programme' should hear this as marking out the conventions that suffice for an audience to recognize the primary communicative intentions of a producer of the sentence: as marking out the conventions that are up to the kind of job that Grice sees as crucial.[14]

We can now say why examples like 'Boo' and 'Hurrah' are of no use to expressivists. If these exclamations do not report (in the broad sense operative in this paper) that the speaker has the attitudes they express, that can only be because the conventions whereby they are linked to those attitudes are not sufficiently entrenched or clear-cut; the exclamations may suggest the presence of suitable attitudes, in the way in which the use of 'but' suggests a sense of contrast between the matters it conjoins, but they do not report those attitudes. It follows that if expressivists are to use the exclamation analogue, they have to take the view that ethical sentences serve only in a very loose fashion to signal the presence of the attitudes they express. And that means trouble.

Expressivism is not the view that there is only a rough, unclear, insufficiently-entrenched-for-truth-aptness connection between ethical words and the attitudes that lead to their production. Expressivists think that ethical terms are well-suited for the task of discussing the attitudes they hold ethical sentences express; if the ethical terms were not well-suited, expressivists would have trouble telling us what their books and articles are about. But this means that they cannot say that the reason ethical sentences lack truth conditions is the same as the reason 'Boo' does, namely, that we have not settled the usages precisely enough.

You might, of course, advance a version of expressivism that was a kind of subjectivism-except-that-we-speak-loosely view. This would be an expressivism that holds that all that stops subjectivism being true is that 'right' and the other ethical terms are rather like 'but', and it would secure the doctrine that ethical sentences lack truth conditions. But then the very

[14] For detailed accounts, see Grice, 'Meaning', *Philosophical Review*, 66 (1957): 377–88, Lewis, *Convention*, Cambridge, Mass.: Harvard University Press, 1969, and Bennett, *Linguistic Behaviour*, Cambridge: Cambridge University Press, 1976.

moment that you make suitably precise in words the attitudes you wish to place at the centre of your account of ethics, you will automatically give ethical sentences containing them truth conditions. This highlights the difference between this bizarre style of expressivism, and standard expressivism. Standard expressivists do not hold that once they have made clear in words which attitudes are central in ethical theory on their view, ethical sentences thereby acquire truth conditions. They do not hold that they can only be precise at the cost of refuting themselves.

5. *Third reply: the command analogue*

Some expressivists take as a model of ethical sentences, not exclamations, but rather commands or recommendations.[15] This suggests a further reply that they may make to the argument from Locke. They may try to elucidate the needed sense of expressing – the sense in which ethical sentences express attitudes without thereby having truth conditions – by saying that it is like the relation between the desires and attitudes that lead to an order or a recommendation: they may hold that my 'ethical' pro-attitude to X stands to 'X is right' as my desire that the door is shut stands to the order 'Shut the door'. As orders and recommendations express the desire for certain states of affairs to be realized, and yet apparently lack truth conditions, so they will argue that ethical sentences may have a similar expressive but not truth-conditional role.

The trouble for this suggestion is that it is very plausible that orders have truth conditions in the relevant sense – the same is true of recommendations but we will make the argument in terms of orders. The issue as to whether orders have truth conditions bears on what we are doing when we produce the words 'Shut the door' in the right circumstances (production in a play or a philosophy of language tutorial doesn't count). Every competent English speaker knows that producing the words 'Shut the door' in the right circumstances is *ipso facto* to command that the door be shut; that's what is being done. And that is to say that it makes no difference whether I say 'Shut the door' or I say 'I command that the door be shut' – a point which is independently plausible. But 'I command that the door be shut' obviously has truth conditions. And so it follows that 'Shut the door' has truth conditions too: it is true in S's mouth at *t* just if S did indeed command at *t* that the door be shut.

[15] See, e.g., Stevenson, *Ethics and Language*. Incidentally, it might be argued that the reason 'boo' and 'hurrah' lack truth conditions is not that the conventions governing their use are not 'good' enough, but rather that they belong with commands and recommendations. In this case, we would not have different models for understanding the expressivist view, but the same model approached in slightly different ways.

It might be objected that it is 'crook' English to talk of orders as true or false, but our topic is not what is or is not good English. It might be objected that this fact about English usage points to a deep feature of orders. We would disagree, but the point could in any case hardly help someone who holds that ethical sentences lack truth conditions, for it is good English to talk of them as true or false. It might be objected that 'Shut the door' and 'I command that the door be shut' cannot have the same truth conditions on the grounds that only the latter sentence can appear in certain contexts; for example, as the antecedent of a conditional. This would be an awkward point for expressivists to press for, as everyone agrees, ethical sentences can appear in such contexts. Moreover, there is a plausible explanation for why 'Shut the door' cannot figure as the antecedent of a conditional: its syntax signals that whenever it is used in serious speech, it is being used to *give* an order and it makes no sense for a sentence within the antecedent of a conditional to be used in that way. This means that, in a wide sense of meaning, 'Shut the door' and 'I command that the door be shut' do not mean the same, because they behave differently in certain contexts; but it does not mean that 'Shut the door' in my mouth in 'simple' contexts does not have the same truth conditions as 'I command that the door be shut'

We can also argue for the conclusion that orders have truth conditions from Locke's observations on voluntary signs. Ordering is something we learn to do, and recognize one another as doing, through our shared mastery of the voluntary conventions that settle when we are, and when we are not, ordering. The possibility of this rests on our *knowing* when we are ordering and when not. And so it is plausible that when I use a sentence like 'Shut the door', I believe that I am giving an order and that I am intentionally producing the sentence as a conventional sign of giving an order. Of course, producing the sentence 'Shut the door' – and, equally, the sentence 'I command that the door be shut' – in the right circumstances *is* to order you to shut the door; the sentences are performatives. But, as many have observed against J.L. Austin, this does not stop them having truth conditions; instead, it makes it easy to make them true.[16]

6. *Expressivism and moral disagreement*

Many expressivists with whom we have discussed the problem raised here for their position remind us that philosophy is hard. They insist that

[16] See esp. Lewis, 'General Semantics', reprinted in *Philosophical Papers*, Vol. I, New York: Oxford University Press, 1983. The view that a sentence cannot both be a performative and have truth conditions is prominent in J. L. Austin, *How to do Things with Words*, Oxford: Clarendon Press, 1962, see esp. Lecture I.

because there is an overwhelming argument – the argument from the persistence of moral disagreement – for the view that ethical sentences lack truth conditions, we should simply look harder for the account they need; namely, the account that denies truth conditions to ethical sentences without running afoul of Locke's point and without turning expressivism into subjectivism-except-that-we-speak-loosely.[17]

In this final section, we seek to undermine this motivation by arguing that expressivists are no better placed than anyone else to handle the problem posed by the alleged persistence of moral disagreement. In arguing this, we are arguing against many non-expressivists, as well as against nearly all expressivists. Many non-expressivists concede, grudgingly, that expressivism is especially well-placed to handle the persistence problem.

The argument from the persistence of moral disagreement starts from the observation that two claims are independently plausible. The first is that if you utter the words 'X is right' and I utter the words 'X is wrong', we are in disagreement. It is *not* like my saying that I like X, and your saying that you dislike X. The second is that this disagreement might survive our agreeing about all the relevant facts concerning X, including all facts about psychological responses to X. An example might be the disagreement between us if one of us is a total utilitarian and the other an average utilitarian, in a case where the difference matters for whether or not X is the right thing to do. We might agree about the effects of X on, say, happiness; we might agree that only effects on happiness matter for rightness and wrongness; and we might know all about our difference in attitudes. Still, it seems, we could be in genuine disagreement – it seems, for example, wrong to infer that you and I must be giving the word 'right' a slightly different meaning. But then, it follows, urge expressivists, that 'X is right' and 'X is wrong' do not make claims about the facts, about how things are, for we agree completely about how things are. And if they do not make claims about how things are, they lack truth conditions and must be construed along expressivist lines.

There is a large literature devoted to this argument, and there is no doubting the interest of the issues it raises. Our concern here, though, is simply to point out that it is no kind of argument for the view that ethical sentences lack truth conditions, and no kind of argument for expressivism.

The key question that needs to be addressed is what is meant by *disagreement* in the argument. Often, what is meant when it is observed that if I produce S and you produce T, we are in disagreement is that S and T cannot be true together. This, though, cannot be what expressivists mean, because, for them, *no* ethical sentences can be true together as they cannot be true to start with.

[17] Blackburn has especially urged this on us.

However, if you say 'I believe that Chicago will win the championship' and I say 'I believe that Boston will win the championship', we count as disagreeing, despite the fact that the sentences we produce can be, and most likely are, true together. But if a difference in belief can count as a disagreement, why not a difference in certain sorts of attitude? So, expressivists can fairly count your assertion that X is right and mine that X is wrong as disagreement by virtue of the difference in our moral attitudes. And this, of course, is what they mostly do.[18]

But *if* disagreement in moral attitude counts as disagreement, then the persistence of moral disagreement is no reason to favour expressivism over a subjectivism that holds that our assertions report our moral attitudes. The difference between expressivists and subjectivists over the semantics of ethical sentences is neither here nor there as far as our having different moral attitudes go. Indeed, almost every party to the debate in meta-ethics believes that if I sincerely assert that X is right and you sincerely assert that X is wrong, we must have different moral attitudes; so, if *that* counts as our disagreeing, as expressivists who are not eliminativists about moral disagreement must allow, almost every party to the meta-ethical debate can respond to the problem of moral disagreement simply by noting that a difference in moral attitudes can survive agreement over all the facts.

We think that the result that expressivism does not help with the problem of moral disagreement is no surprise. Moral disagreement, and indeed disagreement in general, is a psychological phenomenon. The production of sentences makes public our disagreements; it does not create them. Your disagreement with Hitler's moral outlook did not come into existence when you remarked on it in language. This means that a thesis whose essential distinction from other theses lies in a certain claim about the semantics of sentences cannot be in a privileged position when it comes to giving a good account of moral disagreement.[19]

Philosophy Program, RSSS,
Australian National University, ACT 0200, Australia
fcj@coombs.anu.edu.au, pnp@coombs.anu.edu.au

[18] Ayer is a notable exception. In *Language, Truth and Logic*, around p. 110, he argues that expressivists should be eliminativists about ethical disagreement. Of course, expressivists who hold that a difference in attitude counts as disagreement, will construe attitude here in a special sense–perhaps as one designed to *spread*, somewhat as belief is designed to fit the facts and desire *simpliciter* is designed to make the facts fit it.

[19] We are indebted to many discussions with many audiences, mention Michael Smith, Martin Davies, Susan Wolf, Peter Railton, and apologize to those whom we should have remembered.