

Divine Hoorays: Some Parallels between Expressivism and Religious Ethics

NICHOLAS UNWIN
Lancaster University

Divine law theories of metaethics claim that moral rightness is grounded in God's commands, wishes and so forth. Expressivist theories, by contrast, claim that to call something morally right is to express our own attitudes, not to report on God's. Ostensibly, such views are incompatible. However, we shall argue that a *rapprochement* is possible and beneficial to both sides. Expressivists need to explain the difference between reporting and expressing an attitude, and to address the Frege-Geach problem. Divine law theorists need to get past the Euthyphro dilemma, and to avoid moral externalism. This paper shows how a combined theory helps us to achieve this.

1. Divine Law Theories and the Euthyphro Dilemma

It is widely held that morality cannot be grounded in religious authority. Even if God exists, and even if divine commands are always right, this (it is often argued) does not establish the relevant kind of authority. The main problem was originally formulated by Plato in the *Euthyphro*, and it takes the form of a dilemma. If we are to equate the following:

- (1) It morally ought to be the case that p
- (2) God commands that p

then we must ask whether (1) is true because of (2), or vice versa. If the former, then God's authority is arbitrary, and therefore morally objectionable; but if the latter, then it becomes irrelevant since the

moral situation has to obtain in advance and independently of it. Either way, it is immediately concluded, God cannot be the source of morality in any useful sense.

Does this really end religious ethics? Many have felt that this is far too swift, and there are two important counter-strategies.¹ Firstly, one may insist that it is only part of ethics, typically the theory of moral obligation, that is required to be grounded in divine commands. As long as some of ethics is allowed to be theologically ungrounded then the remainder will thereby be constrained, and this may answer the charge of arbitrariness. Secondly, one may insist that there is not enough gap in meaning between (1) and (2) for the dilemma to amount to anything in the first place.

The first strategy has some plausibility. The claim that (specifically) moral obligation has little meaning outside its original religious context is endorsed by many, including theists such as G.E.M. Anscombe and P.T. Geach, as well as atheists such as Bernard Williams.² Its drawback is that, unless the part of ethics that is grounded in religion is indispensable, then it is still open to us to reject religious ethics. Williams, for example, argues that the concept of moral obligation is something that we should be better off without, and there is nothing in this first strategy to stop him. The second strategy is more ambitious, and is the one with which we shall be concerned here. It has a certain naturalness in so far as most religious people find the 'dilemma' quite unreal, and we should hesitate before supposing that this can only be because they are philosophically insensitive. Furthermore, if correct, it really would prevent the problem from arising. After all, nobody imagines that Kantian ethics is embarrassed by the inability to decide whether something is wrong because the categorical imperative forbids it, or whether it forbids it because it is antecedently wrong, since there is not enough difference to begin with. Likewise, it is not embarrassed by the thought that torturing children really would become obligatory if the categorical imperative were to command it. The charge of arbitrariness will not stick. However, this strategy is philosophically very controversial if only because it is widely supposed, even by religious people, that the concept of God cannot play such a direct role in the actual *meanings* of ethical judgements. Moreover, God is normally understood to be a free agent in the way in which the categorical imperative, for example,

¹ There are, of course, many divine law theories which accept one or other horn of the Euthyphro dilemma. However, we shall not examine them here. For an excellent summary of such theories and their variations, see Mark Murphy, "Theological Voluntarism." See also Paul Helm (ed.), *Divine Commands and Morality*.

² G.E.M. Anscombe, "Modern Moral Philosophy;" P.T. Geach, "The Moral Law and the Law of God;" Bernard Williams, *Ethics and the Limits of Philosophy*.

is not, which perhaps spoils the above analogy.³ Nevertheless, this paper will argue that this option makes a good deal of sense. The way forward is to look at expressivist metaethics.

2. Expressing versus Reporting Attitudes

Expressivists claim that moral sentences such as ‘It ought to be the case that *p*’ (briefly, ‘*Op*’) lack descriptive meaning and ordinary truth-conditions. They may look as though they report moral facts, but actually they express non-cognitive attitudes of various kinds (e.g., feelings, desires, and so forth). Sentences such as ‘*A* accepts that *Op*’ have truth-conditions, of course, but acceptance in this sense is not a genuine species of belief.⁴ Rather, accepting that *Op* is more akin to desiring or demanding that *p*. Emotivism and prescriptivism are species of expressivism (broadly construed), and more recent, sophisticated versions have been defended by Simon Blackburn⁵ and Allan Gibbard.⁶ Such views are controversial and have well known difficulties, notably the ‘Frege-Geach problem’ of explaining how non-truth-apt sentences such as ‘*Op*’ can form parts of complex formulae, such as ‘If *Op* then *Oq*,’ and thereby enter into logical relations (if such formulae lack truth-conditions, then we lose the standard way of explaining such matters).⁷ Nevertheless, expressivism seems to follow from three plausible premises: (a) that our moral opinions are essentially connected to motivation (moral internalism); (b) that (genuine) beliefs and affective/conative states (‘passions’) are ‘distinct existences,’ in Hume’s phrase; and (c) that if a sentence is genuinely descriptive and truth-apt, then its acceptance counts as a genuine belief.⁸ For this reason, it is a theory that deserves to be taken very seriously. True, our own view will cast

³ Obviously, there is considerably more to the connection between the categorical imperative and free agency than is indicated here. See, for example, Christine Korsgaard, *The Sources of Normativity*, pp. 90–130.

⁴ The use of ‘accept’ simply as a more neutral term than ‘belief’ is not entirely satisfactory, if only because the acceptance/belief distinction is understood rather differently elsewhere, for example in the philosophy of science. However, it has become standard in the expressivist literature.

⁵ “Attitudes and Contents.”

⁶ *Wise Thoughts, Apt Feelings*.

⁷ Prescriptivists also have this problem, and address it by means of a distinction between phrastics, tropics and neustics. However, although such an approach links up usefully with the Kantian concerns mentioned in the previous section, our preferred approach makes use of fewer basic concepts.

⁸ On this, see Michael Smith, *The Moral Problem*, p. 12 *et passim*. See also Linda Zagzebski, “Emotion and Moral Judgement,” for a version of expressivism where the Humean premise is rejected, an idea which she combines with religious ethics in her *Divine Motivation Theory*.

some doubt on premise (b), but in an unusual way and one which is consistent with a good deal of the Humean outlook. However, expressivism's most immediate implausibility is that it seems to make ethics highly subjective. In this respect, it appears to be at the opposite end of the spectrum from religious ethics. However, we shall see that its subjectivity can be contained, and that the resemblance between expressivism and divine law theories is uncanny.

What does the objectivity of ethics consist in? The most obvious requirement for objectivity in any area is that we should be able to sustain a robust distinction between something's actually being the case and someone's merely thinking that it is. If ethics is grounded in nothing more than subjective attitudes of some kind, then it looks as if this requirement cannot be met. How can I be talking about anything objective if I am merely talking about my own feelings, it may be protested? Indeed, how can you and I even be talking about the same thing when we discuss ethics? The standard response, of course, is that moral judgements are to be understood as *expressions* of our attitudes, not *reports* that we have them. Yet this distinction is elusive and too often taken for granted. After all, how can I express my attitude without also reporting that I have it (and vice versa)? Frank Jackson and Philip Pettit have recently argued that no valid distinction can be drawn here.⁹ However, we shall explain the difference in terms of language games. In what we shall call the 'reporting game,' if *A* says 'I have attitude φ ' and *B* says 'That's true' (or 'I agree'), then *B* means that *A* has attitude φ . By contrast, in what we shall call the 'expressing game,' *B* would mean that he himself has attitude φ . Thus disagreements that arise when playing the expressing game reflect clashes of attitude, rather than factual disagreements over whether someone actually has the attitude in question.¹⁰

This idea sounds rather odd, and it is open to an immediate objection, namely that meanings are being wilfully perverted. If *A* says 'I have attitude φ ,' then, it may be protested, that can *only* mean that *A* has φ ; to suppose that what is being said is that *B* has φ is just absurd. If this is doubted, then consider what happens if we attempt to play the expressing game with other kinds of self-ascriptions. For example, a dialogue of the form: "I'm wearing a red shirt" / "That's not true, I'm wearing a blue shirt" / "Look, are you blind? My shirt is *red!*," and so on, is simply insane; and we might wonder why the game will be any saner when played with attitude self-ascriptions. This is indeed

⁹ "A Problem for Expressivism."

¹⁰ This is an idea I introduced in "Can Emotivism Sustain a Social Ethics?" Many ideas from that article are used here, especially in this section.

an important point, but it does not undermine our language game analysis of the reporting/expressing distinction. Rather, it simply draws our attention to the fact that this distinction, however we wish to present it, can only be applied to some types of states and not others: for example, psychological attitudes rather than sartorial states. I can report to you what colour shirt I am wearing, but I cannot ‘express’ this state in any relevant way. (I can, perhaps, indicate that my shirt colour is the one that you, and everyone else, *ought* to have; but it is now not the colour itself, but my attitude towards it that is being expressed.) The absurdity of the sartorial expressing game is, indeed, a good way of making this point clear. However, although ridiculous, we nevertheless know exactly what the game requires us to do, which is the crucial point. It does not presuppose a prior grasp of what it is to ‘express’ something.

Games of this type may still sound highly irregular even in principle, but they make perfectly good sense when applied to belief-states. In a dialogue of the form, “I believe that there is a tree in the quad” / “No, I must contradict you on this, I believe that there is *not* one there,” and so on, what has happened is that the prefix ‘I believe that’ has lost its original psychological meaning and has become merely ‘parenthetical’ in Urmson’s sense.¹¹ It amounts to little more than diffident throat-clearing, and does not add to the meanings of the sentences it prefixes. In actuality, what we now have is a reporting game, but concerning the sentence, ‘There is a tree in the quad.’ More generally, the move from an expressing to a reporting game amounts to a shift in attention from ‘I believe that *p*’ to ‘*p*.’ In particular, to play the expressing game with moral beliefs (imagining, for the moment, that there are such things) is, in effect, to discuss moral facts using normal reporting rules—which is what we expect moral discourse to look like. The interesting question is: what happens if we suppose that moral so-called ‘beliefs’ are not genuine beliefs? On the Blackburn view, sentences such as ‘*A* accepts that *Op*’ are to be analysed as ‘*A* hoorays that *p*,’ where ‘hooraying’ is understood to be a (yet to be fully specified) non-cognitive attitude. Ostensibly, matters have not changed much, since ‘I accept that’ can be used parenthetically just as well as can ‘I believe that.’ The trouble is that there is no straightforward sentence which relates to ‘*A* hoorays that *p*’ as ‘There is a tree in the quad’ relates to ‘*A* believes/accepts that there is a tree in the quad,’ for hooraying is understood to be a unitary attitude in the sense that its doxastic and deontic components cannot be separated. Blackburn uses formulae such as ‘*H!p*’ (‘Hooray to *p*!’)

¹¹ J.O. Urmson, “Parenthetical Verbs.”

here, but the whole point is that such formulae are not normal, descriptive sentences with truth-conditions, and the move from

(3) *A* hoorays that *p*

to

(4) *A* accepts that *H!p*

which is the first move in his ‘quasi-realist’ project, is therefore tendentious. It is, indeed, striking that Blackburn says very little about this first stage of this analysis, preferring to devote most of his attention to explaining the meanings of complex formulae, such as ‘*H!p* → *H!q*.’ The meaning of ‘*H!p*’ is somehow determined by the meaning of ‘*A* hoorays that *p*.’ The latter is a wholly descriptive sentence, and therefore not relevantly problematic; but the move from it to ‘*H!p*’ is.¹²

Now, our expressing game approach provides us with the missing piece of the analysis. To extract an independent, non-descriptive meaning of some kind for ‘*H!p*’ from sentences such as ‘*A* accepts that *H!p*,’ it is enough to play the expressing game with them. Here, the game cannot be reduced to a reporting game with a different subject matter, as happened with the tree-in-the-quad example. Rather, it is to be taken seriously in its own right. However, the upshot is that we start to talk rather like moral realists, even though we are not required to take on board the metaphysical baggage that comes with moral realism. In short, we are becoming ‘quasi-realists,’ in Blackburn’s sense.

This account, although unusual, has considerable merits. Firstly, it does justice to Jackson’s and Pettit’s insistence that we cannot express attitudes without reporting them. In a way, we agree, since exactly the same sentences are involved, regardless of whether they are reported or expressed. Yet at the same time, we can see how a semantic difference can emerge; for although we start with the same sentences, we use them differently. We do not need to swallow Wittgenstein whole in order to agree that use shapes meaning. Moreover, our account allows for a wide variety of psychological states to be expressed, which other defences do not automatically do. For example, James Dreier criticizes Jackson and Pettit on the grounds that, on their view, ordinary imperatives (which are expressions of desire) would lack a distinct and characteristic sort of meaning.¹³ This is a fair point, but imperatives are a

¹² For a fuller critique of Blackburn’s proposed solution, see my “*Quasi-Realism, Negation and the Frege-Geach Problem*.”

¹³ “Lockean and logical truth conditions.”

special case since they already play a well-entrenched role in our language. An expressivist requires a rich variety of affective and conative states to be expressed, and for their expressions to have meanings as distinct as these states themselves; and even though there are no currently recognized types of speech-act which fulfil this purpose. Our ethical thought demands and deserves such diversity. A more systematic theory for transforming reports into expressions is therefore needed, one which retains the many distinctions to be found within our attitudes. It just remains to be shown that such expressive language games can really work.

It may seem clear that they cannot. Even if they are not always quite as demented as the sartorial game, it appears at first sight extraordinary that anyone should wish to enter into discourses of this kind—unless they are known to be reporting games in disguise. However, this is not obvious. If we recognize the need for something like moral discussion (to improve social cohesion, for example), and also recognize that, tragically or otherwise, there are no moral facts ‘out there’ which could guide our attitudes in the right direction, then we may have to invent morality by constructing and participating in a certain kind of attitudinal discourse. This would deliberately mimic ordinary factual discourse in order to add coherence and discipline to our attitudes in a way that is similar to the way in which scientific discourse adds coherence and discipline to our factual beliefs. Even if there are no moral facts, we surely still need moral discourse, and our expressing game is exactly that. Of course, it may be insisted that expressivism is a doomed outlook anyway, and simply because it is impossible to have moral discourse without moral beliefs and moral facts to anchor it. This is evidently what Jackson and Pettit think. However, if we allow expressivist moral discourse to be possible at all, we do not add any *further* problems by formulating it in our terms. And we have already seen that there are powerful arguments in favour of expressivism.

The important difference between reporting and expressing games is that, in the former, the pronoun ‘I’ has a fixed reference, whereas in the expressing game, it flits from speaker to speaker rather like the transcendental ego. Let ‘I’ in the latter sense be written ‘î’ (and pronounced ‘i-hat’). The interesting question is how much difference in meaning there can be between ‘I hooray that *p*’ and ‘î hooray that *p*,’ and whether it is enough to simulate the difference between ‘I accept that *Op*’ and ‘*Op*.’ At first sight, there appears to be no real difference at all. However, it should be remembered that two sentences are synonymous only if the biconditional formed from them is a tautology, and

$$(5) \text{ I hooray that } p \leftrightarrow \hat{\text{i}} \text{ hooray that } p$$

is no such thing. For when you consider my utterance of (5), you are considering whether to hooray whatever I hooray. Logic does not demand that you do that, nor does it demand that I think otherwise. Thus neither of us can regard it as a tautology. Of course, I myself cannot assent to one of the equivalents without assenting to the other, but this merely reflects that I cannot assent to only one of 'I accept that *Op*' and '*Op*.' Such discrimination is pragmatically self-defeating even though there is no logical implication in either direction. This, however, is a general point (it is G.E. Moore's celebrated 'paradox of infallibility'), and does not reflect a particular difficulty about expressive meaning. It applies regardless of whatever kind of sentence we might substitute for '*Op*.' True, the word 'believe' in 'I believe that *p*' has, in a way, exactly the same meaning as it does in 'î believe that *p*:' we are certainly not dealing with an orthographic accident of some kind. Yet on the other hand, we still have a genuine difference between the psychological (reporting) and parenthetical (expressive) uses of the word, and this difference is, in a broad sense, semantic.

What do come strangely close to being logical equivalences are the following:

(6) 'î believe that *p*' \equiv '*p*'

(7) 'î hooray that *p*' \equiv '*Op*'

In both cases, accepting only one of each pair is more than just pragmatically self-defeating. We can thus start to see the relevance of religious ethics, since the expressive 'î' is developing an eerie resemblance to the name 'God.'

So, am î God? To investigate this extraordinary question, we must see how closely the sentence-pairs in (6) and (7) really do come to being logical equivalences. At first sight, they do not come very close at all. However, we can only say that two sentences are logically non-equivalent if we can describe a logically possible situation where only one of them is true. This is easy enough when we use the ordinary pronoun 'I.' In the case of (6), I merely have to envisage a situation where my beliefs about *p* are different to what they currently are; in the case of (7), the situation to be envisaged is one where my attitude towards *p* differs from the one I have here and now. A hypothetical conversation between my real self and my counterfactual self would consist of disagreements of various kinds. But can we distinguish between the real î and a counterfactual î in the required way? The problem is that the expressive 'î' flits from context to context in a way that undermines the possibility of serious inter-contextual divergence. For example, if my

counterfactual self says, ‘I believe that not- p ,’ I would ordinarily agree with him, and say that he speaks truly—i.e., that he really does believe that not- p (unless I think that he is insincere, or has somehow misidentified his own beliefs). However, if he plays the expressing game, and says, ‘ $\hat{\mathbf{I}}$ believe that not- p ,’ then I cannot agree with him in this way; for in considering that sentence, ‘ $\hat{\mathbf{I}}$ ’ now refers to me *here and now*, and not how I would be in the counterfactual situation. I have thus not succeeded in envisaging a situation where the right-equivalent of (6) is true and the left-equivalent false. A similar argument undermines any attempt to envisage the latter true and former false. It really is beginning to look, therefore, as though I must treat (6) as a genuine logical equivalence; likewise (7).

Yet there are clearly limits to how far this line of argument can go. We cannot even transform the equivalence in (6) into a single sentence:

$$(8) \hat{\mathbf{I}} \text{ believe that } p. \leftrightarrow p$$

or equivalently:

$$(9) \hat{\mathbf{I}} \text{ have the property } \lambda x[x \text{ believes that } p. \leftrightarrow p]$$

for this states that $\hat{\mathbf{I}}$ am omniscient.¹⁴ Indeed, it could be strengthened into the claim that $\hat{\mathbf{I}}$ am necessarily omniscient (certainly *de dicto*; possibly also *de re*). Yet there is no way in which any human speaker could rationally assent even to the unmodalized sentence, ‘I am omniscient’—and, crucially, this also applies when playing the expressing game. A similar argument with (7) would apparently guarantee

$$(10) \hat{\mathbf{I}} \text{ hooray that } p. \leftrightarrow Op$$

and hence

$$(11) \hat{\mathbf{I}} \text{ have the property } \lambda x [x \text{ hoorays that } p. \leftrightarrow Op]$$

This states that $\hat{\mathbf{I}}$ am morally perfect, which is equally unacceptable, even if we can see that something not unlike it makes a certain sense (as can be seen if we substitute ‘ $H!p$ ’ for ‘ Op ’).

The strange behaviour of $\hat{\mathbf{I}}$ -sentences suggests that they cannot be exactly equivalent to $\hat{\mathbf{I}}$ -free sentences, as the above argument shows. This is not entirely surprising, of course, for the whole point about attitude-expressions is that they are not ordinary truth-apt sentences, and

¹⁴ We use Church’s λ -notation to indicate predicate-abstraction.

the Frege-Geach problem warns us that odd things are likely to happen when we attempt to embed them within larger contexts. Nevertheless, there is a near-equivalence that is sufficiently strong to command our attention. Moreover, it works slightly differently between (6) and (7), as we can see if we attempt to force a Euthyphro-style dilemma by asking whether, in each case, the left-hand side (LHS) is true because of the right-hand side (RHS), or vice versa. In each case, this gives us three possibilities:

- (a) The LHS is true because of the RHS;
- (b) The RHS is true because of the LHS;
- (c) The LHS and RHS are too closely related for a contrast to emerge.

In case (6), the obvious choice is (a). Only a projectivist Berkeley would opt for (b), and only in some cases and because of a particular problem (how to explain what constitutes physical existence in the absence of material substance). Option (c), likewise, is not a serious option, if only because (idealism aside) sentences such as '*p*' have a clear meaning that can be explicated independently of what anybody thinks about them. Moreover, and crucially, such explication not only fails to require an understanding of parenthetical prefixes: it also fails to yield it. In case (7), however, (a) is not a serious option, and precisely because (according to expressivism) the RHS cannot be understood independently of our attitudes. Option (b) looks more plausible, since without attitudes there would be no obligations, on this view. However, (c) is even more plausible, since the whole point is that obligations are not so much created by attitudes as actually constituted by them. The obligatoriness of *p* simply has no attitude-free existence, which means that RHS cannot distance itself from the LHS sufficiently to give us a contrast. This is not entirely convincing, of course, for it still seems as though the LHS is very much the senior partner of the equivalence—if only because it is clearly the analysans, whereas the RHS is the analysandum. However, I suggest that we be satisfied, for the moment, with the answer, 'somewhere between (b) and (c), but rather closer to (c).' What is significant is that, as we noted in §1, this is exactly the answer that it is most plausible to give with

(12) 'God commands that *p*' \equiv '*Op*'

and for interestingly similar reasons.

Still, it may be protested that expressivist theories do not genuinely require any kind of religious support. However, even that is not entirely obvious. Blackburn's version talks of 'hooraying,' but that is just a place-holder for more specific attitudes. At some stage, we shall need to go into more detail, and there may be difficulties here. A well known criticism of emotivism, for example, is that moral judgements do not consist of the expressions of just any kind of emotion. Only specific kinds are involved, and the suspicion is that we cannot identify which kinds without importing ideas which go well beyond the minimal framework with which emotivism presents us.¹⁵ We instinctively know (roughly) what sorts of attitudes are likely to be relevant, and these are the ones that we express and try to get others to share. However, our concept of (specifically) moral obligation has a religious origin, and the attitudes expressed in moral judgements are essentially those originally attributed to God. The suggestion that we can rid ourselves completely of religious concepts and yet still know which sorts of attitudes constitute moral judgements is not obviously right. Gibbard's theory applies to all kinds of normative judgements, and he isolates the moral kind by appealing to a variety of distinctive notions, notably guilt.¹⁶ Although Gibbard's theory is more carefully worked out than Blackburn's at this point, and considerable effort is made to ensure a naturalistic formulation of all key concepts, we can see that religious ideas are still hovering in the background (if only as historical influences).

Moreover, we still have the fundamental problem of explaining just why our moral expressing game should be expected to work at all—why it should be any less demented than its sartorial cousin, for example, or (if the latter is too obviously silly to be worth considering) games where other psychological states are expressed, such as gastronomical preferences. Why should attitudinal convergence be expected in the first place? If human beings are understood to be the creation of a morally perfect being whose attitudes are reflected in those of His creatures, then perhaps an answer can be given. Of course, it is not the only possible answer.

Despite this, the suggestion that expressivists might find it worth their while to 'get religion' in order to help them with their distinction between reporting and expressing an attitude still sounds risible. However, we are not quite committed to saying that! The claim, rather, is that they can usefully examine how religious concepts are sometimes used here. This is not too controversial inasmuch as moral beliefs are

¹⁵ On this, see, for example, Alasdair MacIntyre, *After Virtue: A Study in Moral Theory*, p. 12.

¹⁶ *Wise Thoughts, Apt Feelings*, pp. 297 ff.

sometimes analysed, for example, as beliefs about what a morally perfect being, or the Ideal Observer (IO), *would* want (never mind whether there are such beings). Since moral beliefs are thus understood to be beliefs *about* desires, we can see why they might resemble desires themselves sufficiently well to be able to motivate us in the way demanded by moral internalism.¹⁷ Of course, many problems remain, notably that of how such hypothetical desires (or our beliefs about them) can be guaranteed to influence our actual ones. The move from (7) to (11), and the question of how divine hoorays relate to human ones, is very much to the point in this respect. There is also the risk that the IO can only be morally relevant because the word ‘ideal’ is itself morally loaded: in which case sentences using ‘IO’ are already problematic. What is significant, however, is that theories of this kind cannot sensibly be rejected just because they import religious ideas, if only because the degree of religiosity required is fairly small—at least, at the outset. Moreover, the use of ‘I’ simply as a formal device to help anchor moral discourse has considerable merits as it stands, as we have seen. To see why a greater religious influence might be desirable, we need to approach the matter from the other direction.

3. Divine versus Human Attitudes

We firstly need to sharpen our conception of religious ethics. The term ‘divine command theory’ is normally used here; but, following Philip L. Quinn, we shall sometimes use the term ‘metaethical theological voluntarism’ (briefly, ‘voluntarism’) instead.¹⁸ The reason is that the term ‘command’ is very restrictive. Some voluntarists prefer to speak of what God wills, and there is a lively debate on which version is preferable.¹⁹ Although we shall not enter into this debate directly, it turns out to be important that we consider a wide range of alternatives. The prefix ‘metaethical’ is crucial since it indicates that the theory is a thesis about ethical concepts. By contrast, normative theological voluntarism claims that God’s attitudes have normative force, but does not assume that this automatically follows from the nature of ethics itself.

¹⁷ Smith, for example, although not himself an expressivist, presents a theory not unlike this (in *The Moral Problem*). Gibbard (in *Wise Thoughts, Apt Feelings*, §5) talks about a hypothetical goddess Hera, who plays an important role in his normative logic. Blackburn’s logic of attitudes (in “Attitudes and Contents”) likewise makes heavy use of morally ideal situations. Moreover, some, such as Charles Taliaferro (“The Ideal Observer’s Philosophy of Religion”), have actually identified God with the IO, a position close to ours.

¹⁸ Philip L. Quinn, “An Argument for Divine Command Ethics.” Mark Murphy (“Theological Voluntarism”) also prefers this term.

¹⁹ See, for example, Zagzebski, *Divine Motivation Theory*, pp. 258–70.

What is fundamentally significant here is that both expressivism and voluntarism tend to be concerned with similar types of attitudes. Prescriptivists emphasize universal imperatives, and this is echoed by divine command theories. Likewise, expressivist theories that identify moral thoughts with certain types of desire are echoed in theories that emphasize God's will as opposed to his edicts. By contrast, theories which emphasize affective rather than conative states, such as emotivism, have their analogue in theories which define the good in terms of the objects of God's love. The difference in each case, of course, concerns just who has the attitudes in question, and it might be thought that the differences between God and Man are sufficiently great to ensure that any parallels are of minimal significance. Although we have already noted that the expressive \hat{e} is liable to take on divine characteristics, and for reasons internal to expressivism itself, it may still be insisted that the expressivist's God, like the transcendental ego or the Ideal Observer, can never be anything more than a logical fiction—and therefore very far removed from the God of Abraham, Moses, and Isaiah.²⁰ Yet, not only does expressivism take on a religious tinge when formulated in certain kinds of way, there are, conversely, expressivist implications within voluntarism.

The key point is this. Even if ethics is grounded in divine attitudes of some kind, we still have to explain what it is for a person A to have a given moral belief—that Op , for example. Evidently, she needs to believe that God has the appropriate attitude (call it φ) towards p , and this appears to be a straightforwardly cognitive state, one which reflects the fact that 'God has attitude φ towards p ' is (or seems to be) a wholly descriptive, truth-apt sentence. Yet A is also supposed to be moved towards the bringing about that p if she is to avoid the pitfalls of a purely externalist theory of morality. Someone who says, "Well, I agree that God wants it to be the case that p ; but, frankly, I myself am wholly unconcerned about the matter one way or the other," has a rather peculiar understanding of God and her relationship to the Divine. Indeed, there is surely something 'logically odd' about this combination. The questions, 'Do you really believe that God wants it to be the case that p ?' and 'Do you yourself really want it to be the case that p ?' are ordinarily understood to be intimately related in as much as an affirmative answer to the former typically demands an affirmative answer to the latter. Likewise, if you sincerely think that God commands that p , then, other things being equal, you yourself will issue a universal imperative in the prescriptivist sense, even if you do

²⁰ We shall ignore the non-Abramic religions, since they do not attempt to underlie ethics to anything like the same extent.

not suppose yourself to be the ultimate authority behind the command. Similarly, if you genuinely think that God loves *X*, even though you yourself do not, then you must, at the very least, suppose that this reflects a failing on your part. We are expected to echo divine attitudes; and if someone does not do so, and finds this unexceptionable, then we might naturally conclude that she does not really believe that God has the relevant attitudes to begin with. Moral judgements should not lose their ‘action-guiding force’ just because they are perceived to emanate from God: quite the reverse. Of course, we can fiddle with the details here, and there needs to be room for some well known discrepancies, such as Milton’s Satan’s “Evil, be thou my good,” *akrasía* and so forth. Such cases are exceptional and essentially pathological, however, and regardless of whether we are dealing with religious ethics. They should therefore not be allowed to dominate the proceedings. Even if the connection between our moral beliefs and our non-cognitive attitudes is rather more complex than we have just indicated, there surely needs to be an internal connection of *some* kind; and we have already allowed that moral attitudes form only a proper subset of our affective and/or conative attitudes.

Nevertheless, there is a familiar puzzle here. If our own non-cognitive attitudes are needed in order to energize our moral thoughts, then the risk is that our beliefs about God are going to get pushed out of the picture. As it stands, our moral thoughts have apparently taken on a hybrid character. My thought that wanton violence is wrong, for example, now seems to be a fusion of my belief that God hates wanton violence and my own hatred of it. The two components apparently cannot be assimilated, let alone identified, if only because one is a genuine factual belief and the other is an emotion. Yet our moral thoughts simply do not feel as though they have this composite nature. Rather, God’s hatred and my own seem as one—or, at the very least, intimately connected.

Now, it might be thought that the connection is obvious: I hate wanton violence *because* I believe that God does. Yet if this connection is meant to supply a reason for my attitude, then it is unsatisfactory.²¹ Unless I think that God is arbitrary, in which case we are impaled on the first horn of the Euthyphro dilemma, then I must suppose that God hates wanton violence because it has some bad-making characteristic *C* (it causes unnecessary suffering, for example). He does not merely hate it because He hates it. If I am to share God’s attitude in

²¹ If it is instead understood as a causal claim, then it is less objectionable in this respect, though it becomes controversial in other ways, of course, and of limited relevance.

any meaningful sense, then I must also share His reason for that attitude, which means that I too must hate wanton violence because it has that same characteristic *C*. But we are now impaled on the second horn of the dilemma, for God's attitude is no longer part of my reasons. Perhaps I could share this attitude without knowing what God's reasons are. However, there is an internal connection between attitude and reason: indeed, if *C* and *C'* are different considerations, then φ -ing *X* because it has *C* and φ -ing *X* because it has *C'* are, strictly speaking, different attitudes. This ensures that it is impossible to have the attitude without knowing the reason. Maybe I could be said to know that God has the attitude in question because some religious authority has told me (and does not encourage me to ask questions), but the sort of 'knowledge' in question is of a rather feeble kind, and a seriously dismal picture of human moral understanding is being suggested here.

Of course, many secularists will insist that religious ethics *is* very dismal in this sense, and that if the source of morality is understood to be some transcendent authority, then human autonomy and rationality will inevitably suffer, and with them our whole ethical lives. However, this is to fail to do justice to the full spectrum of religious thought. The Abrahamic religions insist that God is both transcendent *and* immanent. Exactly how He can be both is conceded to be a mystery, and there is disagreement as to the balance, but as long as God is immanent to some extent, then the fundamental duality between divine attitudes and ours that perplexes us can, perhaps, be resolved.

A biological analogy is sometimes used here. If we think of the relationship between God and human beings as being akin to that between a person and the individual cells that constitute him, then the person, considered as a unitary organism, can be regarded (from the cells' point of view) as both transcendent and immanent. Cells are free, autonomous entities in the (admittedly, rather strained and limited) sense they are subject only to ordinary cellular laws. They do not need an extra-cellular authority called The Organism to guide their normal activities, nor would their repertoire be much improved by the introduction of such an august being. The organism is, to that extent, immanent. Yet although the organism is nothing over and above the cells that constitute it (the former is supervenient on the latter), it is equally true that the cells are nothing over and above the organism that they constitute. Attempts to prioritize one over the other are quite pointless. However, despite this, the organism has qualities that cannot be reduced to cellular properties, and so an element of transcendence is involved. In a similar way, the theses, 'Society has no reality apart from the individuals that compose it,' and 'Individuals have no reality apart from the societies which they compose,' are quite compatible with

each other, and many disputes between individualists and collectivists are, for that reason, utterly confused. This remains so even though it is also true that societal properties cannot all be reduced to properties of individuals, and that holistic, social descriptions and explanations of individual human behaviour are often indispensable.

Intercellular communication can thus be imagined, if you will, as a sort of expressing game played by the individual cells. Each movement in the game (e.g., the firing of a neurotransmitter) is no more than an unconstrained expression of the cell's individual nature; and yet it is also an expression of the 'will' of the organism as a whole—assuming, of course, that the cell is not damaged, cancerous or otherwise dysfunctional. The game cannot be properly understood until both these aspects are apprehended and seen as complementary. Similarly, if our own individual moral attitudes are thought of as God's attitudes acting 'through' us, albeit imperfectly, then expressivism and divine command theories can likewise be seen as complementary.

It may be feared that this conception of God is too close to pantheism or panentheism to be orthodox, and that divine transcendence has been downplayed far too much.²² Indeed, it may be protested that the actual religious element implicit in this sort of expressivism is just too meagre to be of any great interest to anyone. Nevertheless, the connection between divine authority and human rationality is a delicate one, and reflects an important tension within religious thought itself. What I shall call the 'optimistic view' regards human reason as a gift from God and the way in which divine wisdom is manifested to us.²³ To lack faith in our own faculty of free and independent judgement is thus to lack faith in God and the way in which He has designed us. On the other hand, what I shall call the 'pessimistic view' regards human

²² Though not everyone agrees. For example, the following extract is taken from the 'The Conservative & Masorti Judaism FAQs' website: "However, there is another view of transcendence which rescues us from the classical paradox. In this view, we can compare the relationship between God and the world to the relationship between a person and the individual molecules and cells of their body. A person is much more than the sum of their parts. One notes that the human mind—which possesses life, consciousness, intelligence and free will—transcends the mere matter of which it is made, which has no properties of mind at all. Similarly, God can be said to be to the universe as the mind is to the body's cells. In this view, God is in some way identical to the universe yet at the same time transcends it; the whole is greater than the sum of its parts, not just in magnitude, but in kind." (<http://groups.msn.com/judaismfaqs/isgodonedialecticalwaysofunderstandinggodsnature.msnw?pgmarket=en-us>) Admittedly, we do not have the explicit assimilation of *persons* to individual cells. But since persons are part of the world, the analogy is presumably implicit.

²³ In this context, 'reason' and 'rationality' just mean our faculty of independent judgement, and include more than our purely intellectual abilities.

reason as a gift from the serpent, one which can only lead to our alienation from God's word. Both views have considerable textual support, and they form indispensable, if uneasily combined, elements within Judaism, Christianity and Islam. The optimistic view emphasizes God's immanence in the manner which we have already indicated; whereas the pessimistic view emphasizes His transcendence—the fact that, in our fallen state, we are now alienated from Him. There is not an exact correlation, of course, but nevertheless a useful parallel.²⁴

The convergence between divine law theories and expressivism is most obvious when we stress the optimistic view. Thus imagine a community of basically good people, all made in God's image, but with only finite mentality and subject to earthly temptations. How will they work out what is right and wrong? If they are voluntarists, they will aim to work out what it is that God commands, wills, loves, and so on. To do this, they can only rely on what they themselves command, will, love, and so on. Other sources are possibly corrupt and can only be evaluated using their own powers of judgement. They know that their own attitudes are highly fallible, but have faith that God made them in such a way that His voice may be heard if they search their hearts and their minds in the right way. At any rate, they know that there is nothing available to them that is any more reliable than their own faculties. Such an outlook is massively different from the secularist's in many ways, of course. Nevertheless, their divine metaethics, so formulated, does not merely resemble expressivism in a few curious and unexpected ways. Rather, it *is* a species of expressivism! Human affective and conative attitudes (or, at least, some of them) are characterized here as the internalizations of God's own attitudes, which is philosophically unusual nowadays; but moral judgements are still deemed to be nothing over and above their 'expressions' in the sense understood by expressivists.

If this idea still sounds dangerously close to Enlightenment secularism, consider, for example, the deliberations of cardinals in conclave. When debating who should be elected pope, each cardinal supposes his conscience to be, although not identical with God's voice, at least directed by it. No cardinal would treat his own voice and God's as potentially *competing* authorities. If he is tempted into error, he knows that it is not because his conscience itself is 'off-message,' if the expression may be excused, but because he has ignored or distorted what it is actually telling him. Now, does this reliance on individual conscience

²⁴ A historical complication is that many Enlightenment thinkers were deists (as well as 'optimists' in our sense). One of their motivations was to keep God out of secular life, yet a consequence of deism is that He becomes wholly transcendent.

amount to a betrayal of religious authority? Have the cardinals succumbed to the pride of Lucifer in so far as they rely solely on their own powers of judgement? We are surely not forced to think so; and the more general assimilation of conscience—even the conscience of an unbeliever—to God’s voice has biblical justification.²⁵

It may be protested that we have ignored the role of special revelation, and that it is just not true that the cardinals (or, indeed, anyone else) are required to rely solely on their own powers of judgement. Revelation is perhaps fundamental to religious belief, and it does indeed give rise to special problems. However, the relevant sort of autonomy is surely still required here. Cardinals in conclave are presumably not directed or guided by an outside agency in the same sort of way in which mediums purportedly are at a *séance*. Revelation may influence our ordinary faculties, and in a very powerful way; but it does not actually overwhelm them. At least, it certainly should not.

The resemblance between expressivist and divine law theories concerning the source of conscience may still look unconvincing. After all, the latter suppose that God is a real entity, whereas the expressivist ‘ Υ ’ has no actual reference at all. Yet problems about the name ‘God’ will arise even from within voluntarist theories, as the following argument shows. We begin with two uncontroversial background premises:

(13) If ‘ S ’ is a descriptive, truth-apt sentence, then the thought that S (i.e., the thought with content $\langle S \rangle$) can be a genuine belief (and nothing over and above a belief)

(14) ‘ A φ s that p ’ is a descriptive sentence for any name ‘ A ’

However, we have already agreed that God’s attitudes need to be internalized if they are to have any moral significance for us, and this yields the entailment:

(15) Necessarily, if A sincerely thinks that God φ s that p , then A also φ s that p .²⁶

²⁵ “Indeed when Gentiles, who do not have the law, do by nature things required by the law... they show that the requirements of the Law are written on their hearts, their consciences also bearing witness, and their thoughts now accusing, now even defending them” (Romans 2:14f). Of course, there is considerable debate as to exactly how conscience and God’s word are related, but we shall not pursue the matter here. At the risk of ignoring too many important issues, we shall simply treat the word ‘conscience’ as the name of our ordinary faculty for moral judgement, and not assume anything further about its nature.

²⁶ The qualifier ‘sincerely’ may need to be amplified to avoid the problems about Milton’s Satan, *akrasia*, and so forth, that we touched on earlier.

However, although voluntarist theories may differ as to what φ should be if it is to ground moral obligation, it is agreed that:

(16) φ is a non-cognitive attitude

That is to say, moral obligation is not grounded in what God believes to be the case, but rather in what He commands, wills, loves and so on. But (15) and (16) directly imply:

(17) ‘ A sincerely thinks that God φ s that p ’ cannot (just) express a belief

for beliefs and non-cognitive attitudes are, by definition, different. But (13) and (17) together yield:

(18) ‘God φ s that p ’ is not a descriptive sentence

This, together with (14), yields:

(19) ‘God’ is not a name

There has been much debate, of course, as to whether ‘God’ really is a proper name, but usually the alternative view is that it is, instead, a title. This is not the problem that we have here. The above argument would apply if we replace ‘name’ by ‘designator,’ understood in the broad Kripkean sense which embraces both proper names, titles, definite descriptions and any other referential (or pseudo-referential) terms. What we have, rather, is a more remarkable argument that suggests that God cannot be referred to, that He has a fundamentally elusive quality not unlike the expressive $\hat{\mathbf{i}}$. Only this time, we reached this conclusion from the other direction, starting from a religious perspective.

It may be suspected that we have smuggled in expressivist assumptions, or at least the premises of our original argument that apparently shows that expressivism is the only viable option. Premise (16) is especially questionable here, for although theological voluntarists may agree that if ‘God φ s that p ’ is to be a plausible analysis of ‘It morally ought to be the case that p ,’ then φ must be a conative and/or affective attitude of some kind, the conclusion that such attitudes must therefore be non-cognitive imports a crucial Humean assumption that, it may be felt, they are not forced to accept. Or if they are, it is only because everyone is forced to accept it, and for reasons that have nothing to do with the religious framework within which the argument has been cast. However, this can also be disputed. The standard modern argument

for Hume's thesis that beliefs and 'passions' (i.e., affective and conative states) are 'distinct existences' hinges on the different 'directions of fit' possessed by these types of attitude. Beliefs have a 'world-to-mind' direction of fit, which means that, should there be a mismatch between the world and how we believe it to be, then the mistake is in our belief, not the world. By contrast, desires, for example, have a 'mind-to-world' direction of fit, which means that, should there be a mismatch between the world and how we desire it to be, then the mistake is in the world, not the belief. Beliefs aim at truth, whereas desires aim at satisfaction. It seems to follow, therefore, that no state can have both directions of fit, for we are then given inconsistent instructions on how to repair a mind-world mismatch.

This argument is widely accepted, but it contains gaps. Most notably, Margaret Little has pointed out that, at best, it only proves that a given state cannot have both directions of fit with respect to the same content; and we are not required to suppose this.²⁷ Nobody needs to think that, for some p , the belief that p = the desire that p . Rather, what a non-Humean supposes is (roughly) that the belief that Op = the desire that p , and $\langle Op \rangle$ and $\langle p \rangle$ are plainly different contents. There is no obvious reason why a single state should not have a world-to-mind direction of fit with respect to the former and a mind-to-world direction of fit with respect to the latter. No inconsistent instructions are involved.

This is an important general criticism; but the crucial point is that the possibilities that it admits are not available to the sort of voluntarist theories with which we are concerned. This is because they require that God's belief that Op should not only be equivalent to (e.g.,) His desire that p ; it must be *constituted* by it, or else ' Op ' remains unanalysed. If an independent analysis of ' Op ' could be given, then we must ask why we still have anything like a divine command theory in any useful metaethical sense. As we have already noted, unless we are to be impaled on the second horn of the Euthyphro dilemma, moral reality cannot be allowed to be separated at any level from divine attitudes. By contrast, a moral cognitivist internalist who follows Plato—such as Little herself, for example—need have no problem with a dual direction of fit here, for she is not required to suppose that the content of her belief that Op is actually made true by her, or anyone else's, desire that p . Quite the reverse. However, such an option is not available to God if He is to be the *author* of moral obligation in the strong sense required here, and not merely a perfect cognitive tracker of

²⁷ Margaret Olivia Little, "Virtue as Knowledge: Objections from the Philosophy of Mind."

antecedently existing moral truth. If His desires actually constitute the obligation, then they cannot, even *qua* beliefs, have a world-to-mind direction of fit. The very idea that God might need to adjust or ‘fit’ His beliefs so as to match moral reality is to ensure that moral reality has an unacceptable degree of independence. The upshot is that divine commands cannot be thought to have the independently formulable credal content that is needed to block our argument at stage (16). Thus points (13) to (19) really do yield an independent argument.

This is not wholly surprising, for the basic difficulty with treating ‘God’ as the name of a distinct entity, from a moral point of view, is already implicit in the Euthyphro dilemma, and has nothing to do with expressivism. Once any such entity has become a person ‘out there,’ identifiable as someone independent of our attitudes, we immediately ask whether His commands are really right; and G.E. Moore’s open question argument ensures that the matter can be disputed. Being ‘out there,’ or ‘other,’ automatically disqualifies Him from moral authorship as far as a truly independent-minded person is concerned. By contrast, as long as God is understood only within ourselves, and knowable only through ourselves, then this objection can, perhaps, be blocked. There is not enough of the right sort of gap between evaluator and evaluatee for there to be an open question of the relevant kind. Nevertheless, as we have just seen, there is a heavy price to be paid if we adopt this strategy, namely that it becomes hard to conceive of God as any kind of distinct entity—‘distinct’ both in the sense of ‘definite,’ and also in the sense of ‘different from us.’

It may be insisted, once again, that the real problem with our account is that we adhere too strongly to what we call the ‘optimistic view,’ for it is this which encourages the unlimited independent-mindedness that underlies the open question argument. If, instead, we emphasize God’s attitudes as *corrective* to human attitudes—and therefore external to them—then, it may be thought, not only will we be able to avoid these technical difficulties about reference; we shall also obtain a view about the relationship between God and humanity that is far more consonant with orthodox religious opinion. However, oddly enough, not only does such an alternative viewpoint fail to solve the problems that originally led us towards the ‘optimistic view;’ it also, and even more surprisingly, fails to distance religious ethics from expressivism.

We need to remember, once again, that only a proper subset of our attitudes are morally relevant, and it is a moot point exactly how we are to identify this subset. A secular liberal might regard a certain kind of hyper-critical emotion as merely neurotic, perhaps a symptom of an overactive superego or something along these lines. By contrast, a

religious conservative might regard it as the real voice of God. The former might try to exclude it from ethical deliberation, whereas the latter might try to exclude almost everything else. The secular/religious and liberal/conservative distinctions do not automatically coincide, of course, and there are many possible combinations of opinion here. The point, however, is that the distinction between what we call 'optimistic' and 'pessimistic' views of human nature cannot amount to anything other than a difference of opinion within this area. Of course, moral attitudes are corrective to human attitudes! We must regard them as such regardless of whether we regard them as God's attitudes. However, if we say that *all* human attitudes are corrupt, and that the voice of God is *wholly* separate from any of our own attitudes (including our own hatred of what are perceived to be our corrupt attitudes), then we end up with a very strange picture, one which no religious pessimist would easily recognize. Even the most extreme Calvinist, for example, would allow that self-hatred of the above kind should be excluded from such blanket condemnation! If our fall from grace were total, if every divine spark within us were extinguished, then we would become not opposed, but simply indifferent to (what we perceive to be) God's commands. Religious ethics would then face the problem faced by all externalist theories, namely that moral judgements would lack any intrinsic motivational force one way or the other. This is a position which all metaethical theological voluntarist theories surely need to avoid, and usually fail to avoid. A person who consistently regarded all his attitudes as wholly separate from God's would be utterly forsaken. We might go further still, and insist that it is he, and not the independently-minded secular liberal, who is closest to Lucifer and eternal damnation. Perhaps this is to overstate the point! Yet it remains intolerable to suppose that God is both the ultimate source of morality, and also that His voice can only be heard if we reject as serpent-driven our native instincts and faculties *tout court*.

All this would be fine, except that expressivism itself has difficulties that we have already expounded, and it is unclear why theological versions are exempt from them. Most notably, moral judgements are deemed to be non-truth-apt, and this gives us the Frege-Geach problem. Sentences such as 'God hoorays that *p*' may look like ordinary descriptive sentences, but the above argument apparently shows that they cannot be. Can we get past this? It may seem that we cannot; but there is a tentative line of argument that could, perhaps, be of use here.

As we have already noted, orthodox opinion is that God is both transcendent and immanent in a sense that is ultimately mysterious, though which can be illuminated by analogies. We have also examined

at length the difference between an empirical ego (such as you or me) and the transcendental ego \hat{I} . Yet, these comparisons are problematic, if only because the transcendental use of 'God' (as similar to ' \hat{I} ') corresponds to an immanent conception of God as a being that has no reality outside of us (and conversely, to treat 'God' as an ordinary proper name is to treat God as a real entity independent of ourselves). Things seem to be the wrong way round! What has happened is that the terms 'transcendent' and 'transcendental' have become entangled, which can easily happen since they each tend to be used slightly differently in different contexts and at different historical periods. The former typically means 'beyond ordinary entities;' the latter is much less easily defined, but it typically includes 'having no empirical existence.' It is only the former that is properly contrasted with 'immanent'; the latter is better contrasted with 'empirically real.' Now, orthodoxy requires that God be real but non-empirical, and this certainly introduces some strain, especially when we ask how God can be similar both to real but empirical people such as ourselves, and also to the transcendental ego \hat{I} . Perhaps we have an impossible combination, and the concept of God, so understood, is simply incoherent. Yet the distinction that can be drawn between the transcendent and the transcendental ensures that this is not obvious. And crucially, if we are already willing to allow that such a being exists, for whatever reason, then it may be that no further mysteries are involved in seeing how God can be sufficiently real that sentences such 'God hoorays that p ' are factual enough to address the Frege-Geach problem, and yet sufficiently unlike ordinary empirical entities to ensure that His commands cannot be 'externalized' in any damaging sense. At any rate, we can see how evaluative expressing games would be stabilized noticeably if this position were accepted.

Of course, we must be sure that there really are no further mysteries involved. This is not easy, since it is hard even to state all the ways in which God, as traditionally understood, is mysterious, let alone to understand them. Yet there are some further intuitions that we can use.

We have already considered the relationship between an organism and its constituent cells, and we touched on the similar relationship between a society and its constituent individuals. Both analogies are limited, for an individual cell is nowhere near intelligent enough to be human, and a society is nowhere near integrated enough to be God. However, the latter analogy is of obvious ethical relevance since morality is widely, and reasonably, thought to be primarily a social phenomenon. That is not only in the sense that our social relationships are the chief source of our obligations, but also in the sense that our individual consciences are themselves socially conditioned. I might think that my

deepest instincts are entirely my own—which, in a sense, is obviously true—and I may feel able to judge, perhaps negatively, the social norms to which I am supposedly subject—again, also obviously true. Yet, despite this, I am a social being, the product of my own social culture. Remove all acculturation, and not much remains: my capacity for intellectual and moral judgement will certainly have gone. Perhaps, then, the expressivist \hat{i} should better be understood as denoting our society rather than God. After all, the idea that each player of the expressing game should think of herself as a mouthpiece of society is by no means implausible.

Yet there is an important way in which this cannot be right. Our society, as it currently is, is not morally perfect, and no sane individual could suppose otherwise, even though moral norms have their origin in society. To analyse ‘ \hat{i} hooray that p ’ as ‘Society hoorays that p ’ is to confuse genuine normative judgements with judgements of descriptive ethics, that is to say, judgements of what are *believed*, rightly or wrongly, to be about what is right or wrong; and attempts to replace the former by the latter can only undermine the reflective criticism essential to ethics.²⁸ Societies are not sufficiently good to be candidates for \hat{i} . But imagine that the society in which we live is a kind of shadow or projection of a perfect society at which we aim. Such an entity, so understood, would be both real and also something like a referent of ‘ \hat{i} .’

Now, religious orthodoxy supposes that human beings are creatures of God in something like the way in which sociological orthodoxy supposes them to be creatures of the societies that they constitute—only in much stronger sense. Moreover, the dependence is no longer symmetric. That is to say, although societies are the products of individuals as well as the other way round, God is not similarly dependent on His creation: His perfection ensures that He transcends it. Perhaps such a picture cannot survive in a secular and scientific age. The fact remains, though, that *if* we adhere to it (as many do), then we have a very powerful tool for dealing with some central metaethical issues. And even if we do not adhere to it, we can learn much from seeing just how a religious framework can make an important difference here.

4. Conclusion

Our thesis is that there are unexpected parallels between expressivism and voluntarism, and that a convergence between them would be beneficial to both theories. Voluntarists stand to benefit the most since there are fewer new assumptions that need to be adopted; by contrast,

²⁸ For a further development of this point, see my “Can Emotivism Sustain a Social Ethics?”

expressivists (most of whom are atheists) would need to make far more radical changes in their outlook. Yet there are certain things which we do not claim to have achieved.

Firstly, we do not claim to have shown that such a combined theory is true. Too many issues have been left unexplored, for example, the prospects for Platonist cognitivist internalism. It also remains unclear just how much progress has really been made with the Frege-Geach problem: our account points towards a solution, but that is all.

Secondly, we have not even shown that expressivists are forced into becoming voluntarists nor vice versa. Although we have shown that such a combined theory has considerable merit, and that there are not all that many alternative options, we have not attempted to explore all the possibilities here, and many things have just been assumed (for example, that all versions of moral externalism are unacceptable).

Thirdly, we do not even claim that expressivists who move towards voluntarism and voluntarists who move towards expressivism must end up with the same views. On the contrary, the difference between secular and religious outlooks remains very profound: specifically the difference between treating God as a real entity as opposed to a purely formal device to anchor moral discourse. Our claim, rather, is just that such differences will not be (for the most part) *metaethical*. They include many factors that can only be treated at a very abstract philosophical level, of course, but (outside pure theology itself) they are primarily metaphysical questions about the nature of persons rather than the nature of values. They include, in particular, the nature of human autonomy, and how the convergence of moral opinion is best explained—questions that we have only touched on. Nevertheless, although our conclusions are fairly modest, the parallels that we have uncovered are sufficiently striking as to show that expressivists and voluntarists should pay considerably more attention to each other than they usually do.²⁹

References

- Anscombe, G.E.M., (1958) “Modern Moral Philosophy,” *Philosophy* 33, pp. 1–19.
- Blackburn, Simon, (1993) “Attitudes and Contents,” in *Essays in Quasi-Realism* (Oxford: Oxford University Press), pp. 182–97.
- Dreier, James, (2004) “Lockean and logical truth conditions,” *Analysis* 64, pp. 84–91.

²⁹ I am grateful to two referees from *Philosophy and Phenomenological Research* for some very useful criticisms of an earlier version of this article.

- Geach, P.T., (1969) "The Moral Law and the Law of God," in *God and the Soul* (Routledge & Kegan Paul), pp. 117–29.
- Gibbard, Allan, (1990) *Wise Thoughts, Apt Feelings* (Oxford: Clarendon Press).
- Helm Paul (ed), (1981) *Divine Commands and Morality* (Oxford: Oxford University Press).
- Jackson, Frank and Philip, Pettit, (1998) "A Problem for Expressivism," *Analysis* 58, pp. 239–51.
- Korsgaard, Christine, et al., (1969) *The Sources of Normativity* (Cambridge: Cambridge University Press).
- Little, Margaret Olivia, (1997) "Virtue as Knowledge: Objections from the Philosophy of Mind," *Noûs* 31, pp. 59–79.
- MacIntyre, Alasdair, (1981) *After Virtue: A Study in Moral Theory* (London: Duckworth).
- Murphy, Mark, "Theological Voluntarism," *The Stanford Encyclopedia of Philosophy* (Fall 2002 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2002/entries/voluntarism-theological/>> .
- Quinn, Philip L., (1990) "An Argument for Divine Command Ethics," in Michael Beaty (ed.) *Christian Theism and the Problems of Philosophy* (Notre Dame: University of Notre Dame Press), pp. 289–302.
- Smith, Michael, (1994) *The Moral Problem* (Oxford: Blackwell).
- Taliaferro, Charles, (1999) "The Ideal Observer's Philosophy of Religion," *Proceedings of the 20th World Congress of Philosophy* 4, pp. 51–7.
- Unwin, Nicholas, (1990) "Can Emotivism Sustain a Social Ethics?," *Ratio* N.S. 3, pp. 64–81.
- , (1999) "Quasi-Realism, Negation and the Frege-Geach problem," *The Philosophical Quarterly* 49, pp. 337–52.
- Urmson, J.O., (1952) "Parenthetical Verbs", *Mind* 61, pp. 480–96.
- Williams, Bernard, (1985) *Ethics and the Limits of Philosophy* (Cambridge MA: Harvard University Press).
- Zagzebski, Linda, (2003) "Emotion and Moral Judgement," *Philosophy and Phenomenological Research* 66, pp. 104–24.
- , (2004) *Divine Motivation Theory* (Cambridge: Cambridge University Press).