The Philosophical Quarterly, Vol. 49, No. 196 *July 1999*
ISSN 0031–8094

QUASI-REALISM, NEGATION AND THE FREGE–GEACH PROBLEM

BY NICHOLAS UNWIN

I

Expressivist analyses of moral language have been of considerable interest for some time. Examples include Ayer's and Stevenson's emotivism, Hare's prescriptivism, Blackburn's projectivism and Gibbard's norm-expressivism.[1] Such theories are diverse, though they all share the fundamental thesis that the function of a moral judgement is not to express a moral belief that aims at truth (in the same sort of way as an ordinary factual sentence does), but rather to express a non-cognitive attitude, a practical commitment, adherence to a system of norms, or something along similar lines. Thus moral sentences which say that something is right or wrong, good or evil, and so forth, are deemed to have only the outward form of ordinary descriptive sentences. In reality, they are something else in disguise; and the expressivist's task is to tell us what they really are, and how they manage to disguise themselves so well.

A central difficulty here is the 'Frege–Geach problem' of unasserted contexts, also known as the 'embedding problem'.[2] The problem is that sentences that express moral judgements can form parts of semantically

[1] A.J. Ayer, *Language, Truth and Logic* (London: Gollancz, 1936); C. Stevenson, *Ethics and Language* (Yale UP, 1944); R.M. Hare, *The Language of Morals* (Oxford UP, 1952); S. Blackburn, *Spreading the Word* (Oxford: Clarendon Press, 1984), and 'Attitudes and Contents', repr. in his *Essays in Quasi-Realism* (Oxford UP, 1993), pp. 182–97; A. Gibbard, *Wise Choices, Apt Feelings* (Oxford UP, 1990).

[2] See P.T. Geach, 'Assertion', *Philosophical Review*, 74 (1965), pp. 449–65; Hare, 'Meaning and Speech Acts', *Philosophical Review*, 79 (1970), pp. 3–24; Blackburn, 'Realism, *Quasi* or Queasy?', in J. Haldane and C. Wright (eds), *Reality, Representation and Projection* (Oxford UP, 1993), pp. 337–63; G. Schueler, '*Modus Ponens* and Moral Realism', *Ethics*, 98 (1988), pp. 492–500; R. Hale, 'The Compleat Projectivist', *The Philosophical Quarterly*, 36 (1986), pp. 65–84, and 'Can There Be a Logic of Attitudes?', in Haldane and Wright, pp. 337–63.

complex sentences in a way that an expressivist cannot easily explain. For example (from Geach), the sentence 'Telling lies is wrong' has the same meaning regardless of whether it occurs on its own or as the antecedent of 'If telling lies is wrong, then getting your little brother to tell lies is also wrong'. This must be so, since we may derive 'Telling your little brother to tell lies is wrong' from them both by *modus ponens* without any fallacy of equivocation. Yet nothing is *expressed* (in the relevant sense) by 'Telling lies is wrong' when it forms the antecedent of the conditional, since the antecedent is not itself asserted, and so its meaning (regardless of where it occurs) apparently cannot be explained by an expressivist analysis. Analogous problems occur with other types of embedded context.

This is the problem as introduced by Geach, who recommends attention to Frege's distinction between assertion and predication. Semantic complexity is evidently an indispensable feature of moral discourse, and moral argument requires the use of logical laws that certainly look exactly the same as the ordinary variety used in factual contexts. Yet the contexts introduced by ordinary logical operators such as 'and', 'or', 'not', 'if ... then ... ' and the quantifiers, together with predication itself, are normally explicated in terms of the more basic semantic concepts of truth and satisfaction. Although the matter is not beyond dispute, it seems that this option is not available to expressivists, since they deny that moral sentences have truth-conditions – or at least truth-conditions of a standard sort.[3] Various solutions have been put forward to deal with this problem, the most sophisticated and developed of them by Simon Blackburn[4] and Allan Gibbard.[5] Both authors recognize the importance of the problem, and a desire to solve it has been a major influence in shaping their respective overall theories. Leaving Gibbard to another occasion, I shall argue here that the simplest logical context of them all, negation, presents Blackburn's theory with very serious difficulties.

II

Blackburn's ethical theory is a development of Hume's projectivism, and his '*quasi*-realist project' is to show how we can earn the right to treat moral sentences as if they were straightforwardly truth-apt, even though their

[3] See, e.g., P. Horwich, 'Gibbard's Theory of Norms', *Philosophy and Public Affairs*, 22 (1993), pp. 67–78, and 'The Essence of Expressivism', *Analysis*, 54 (1994), pp. 19–20; M. Smith, 'Why Expressivists about Value Should Love Minimalism about Truth', *Analysis*, 54 (1994), pp. 1–12, and 'Minimalism, Truth-Aptitude, and Belief', *Analysis*, 54 (1994), pp. 21–6; J. Dreier, 'Expressivist Embeddings and Minimal Truth', *Philosophical Studies*, 83 (1996), pp. 29–51.

[4] *Spreading the Word* pp. 181–223, and 'Attitudes and Contents'.

[5] *Wise Choices, Apt Feelings* pp. 83–102.

actual function is to express conative attitudes of various kinds, not to describe reality. He has more than one reason for thinking this, but the most important line of argument derives from three theses.[6] The first is Hume's thesis that (genuine) beliefs and desires ('passions') are 'distinct existences'. The second is the moral internalist's thesis that (what we call) our moral beliefs are internally connected to the will. Finally there is the implicit thesis that if a sentence is genuinely descriptive or truth-apt, then its acceptance counts as a genuine belief. The first two theses are certainly controversial, but they lie outside the scope of this paper. The third, by contrast, is not relevantly problematic in itself, but it draws attention to the question of what expressions of non-cognitive attitudes (notably, the 'passions') can mean if they lack truth-conditions and do not mean anything descriptive. Evidently it will depend on the attitude, and here Blackburn does not go into much detail. He speaks of 'hooraying' and 'booing', but these are just place-holders for the full range of moral attitudes. Such generality is desirable in as much as his theory needs a general scope, but we need to remember that sentences such as 'Hooray to your telling the truth!' are hardly transparent. If they are to be more than mere ejaculations of the type posited by the crudest versions of emotivism, then more will need to be said at some stage on what it is to hooray something.

Blackburn presents us with a formal apparatus which begins with two attitude operators, '$H!$' and '$B!$', which stand for 'Hooray!' and 'Boo!' respectively. Originally these were understood to govern gerunds such as 'telling lies', but were later revised to govern instead whole sentences. So modified, they function very much as ordinary deontic operators (such as 'It ought to be the case that ...'), though expressions such as '$H!p$' ('Hooray to p') are not themselves descriptive sentences. The chief task, then, is to explain what such expressions do mean. The *quasi*-realist project is primarily concerned with the meanings of complex formulae such as '$H!p \rightarrow H!q$'; but before analysing semantic complexity we evidently need some grasp of the atomic formulae. Although Blackburn does not make this explicit, the '$H!$' operator is definable contextually by

E. A accepts $H!p \equiv A$ hoorays that p.

Because the word 'hooray' appears on the right-hand side as a verb, rather than as a mere ejaculation, it is susceptible to analysis of an ordinary, familiar kind, and such analysis is evidently crucial to Blackburn's overall theory. For example, the way to explain in detail why his theory is both non-cognitivist and non-descriptivist is by telling an appropriate story about the

[6] See *Spreading the Word* pp. 181–9; also Smith, *The Moral Problem* (Oxford: Blackwell, 1994), *passim.*

truth-conditions of 'A hoorays that p'. The latter sentence is uncontroversially descriptive, and so there are no relevant theoretical difficulties involved in giving it a precise and complete analysis. Such an analysis does not require a prior understanding of the formula '$H!p$', but it will nevertheless ensure that '$H!p$' can mean at most one thing. This is because if two sentences (descriptive or not) are to have different meanings, then it must be logically possible for someone in some circumstances to accept one but not the other, and the (E) schema should be understood as a universally quantified logical equivalence. This equivalence will therefore provide the beginnings of something like a Russell-style contextual definition of the '$H!$' operator (or of the more precise operators for which it is a place-holder). But only the beginnings, for a full-blown semantics must include an account of how it can appear in indefinitely many indirect contexts. This, of course, is exactly the Frege–Geach problem; and it is a problem because it is unclear where the additional information can come from.

Blackburn's solution is to develop a logic of these attitude-operators, and in such a way that the semantics of complex formulae can be explicated through their inferential roles. When dealing with ordinary sentences, we define an argument to be valid if the premises and the negation of the conclusion form an inconsistent set, that is to say, a set whose members cannot all be true at once. However, this is to invoke a conception of truth which is not available to expressions of attitude as presently construed. An analogous notion is required, and Blackburn claims that 'The solution, I believe, must rely upon a notion of consistently realizable attitude. If you promote p and promote $\sim p$, you are in a practical analogue of an inconsistent position, in which not all your goals can be realized.'[7] In general, the notion of consistent realizability (as understood in imperative logics, for example) is claimed to be extendable to expressions of attitude so as to produce a similar kind of logic, one which sufficiently resembles ordinary systems to warrant the name 'logic'. Crudely speaking, the issue is whether everything that is being hoorayed is capable of being simultaneously true. All that remains is to work out the details.

Although Blackburn's solution to the Frege–Geach problem applies to all sentential operators, the behaviour of the binary operators, rather than the negation operator, is his chief object of attention. This is for two reasons. First, Blackburn is concerned primarily with moral inference, especially moral *modus ponens*, and so expressions of the form '$H!p \rightarrow H!q$' need particular attention. Second, binary operators need to be able to combine expressions of attitude with ordinary truth-apt sentences (as in '$p \vee H!q$', for

[7] Blackburn, 'Gibbard on Normative Logic', *Philosophy and Phenomenological Research,* 52 (1992), pp. 947–52, at p. 949.

example), and this is especially difficult to explain if expressive operators are thought to differ in meaning from their truth-functional analogues. Bob Hale in particular has sought to pin Blackburn down on the question of how such operators can both govern (non-truth-apt) expressions of attitude and yet also sustain what appear to be ordinary logical inferences. Consequently they both concentrate on the extent to which conjunction and disjunction (and hence the conditional, construed extensionally) can be captured by Blackburn's preferred formal semantics, namely a tableau system based on Hintikka's system of model sets;[8] and the debate focuses on various details concerning this apparatus.

As far as negated formulae are concerned, however, Blackburn's strategy is simply to introduce a dual operator '$T!$' and rephrase '$\sim H!p$' as '$T!\sim p$', a manœuvre which he finds unproblematic: 'corresponding to permission we can have $T!p$ which is equivalent to not hooraying $\sim p$, that is, not booing p' ('Attitudes and Contents' p. 189). Yet this is just wrong, since 'not hooraying $\sim p$' means 'not accepting $H!\sim p$', as opposed to 'accepting not-$H!\sim p$', which is what permission amounts to. This may seem a small detail, but in fact it points to a major difficulty. In other contexts, there is certainly a very real difference between not accepting (or refusing to accept) something and actually accepting its negation. For example, atheists and agnostics alike refuse to accept that God exists, though only the former accept that he does not. More basically, to accept the negation of a sentence S is to accept *something*, whereas to refuse to accept S is consistent with accepting nothing at all. It may be protested that this distinction involves a realist conception of acceptance-transcendent truth which an expressivist is not required to endorse. But an anti-realism which denies that one may consistently refuse to accept a statement and equally refuse to accept its negation is far too extreme to be acceptable in ethics. For instance, the question of whether Dr X ought to perform surgery on patient Y is one on which I may legitimately refuse to hold any opinion either way. Yet if the acceptance of not-S is understood to be nothing over and above the refusal to accept S, then a two-way refusal both to accept S and to accept not-S will involve both accepting and refusing to accept the same thing (namely not-S), and this is clearly absurd.

Obviously Blackburn's suggestion is wrong as it stands, but it is not easily amended. This whole issue arises because of a fundamental syntactic defect with the original equivalence schema

E. A accepts $H!p \equiv A$ hoorays that p

[8] J. Hintikka, 'Deontic Logic and its Philosophical Morals', in his *Models for Modalities* (Dordrecht: Reidel, 1969), pp. 184–214.

namely, that the left-equivalent admits of three negated forms, whereas the right-equivalent admits of only two:

N1. *A* does not accept $H!p \equiv A$ does not hooray that *p*
N2. *A* accepts not-$H!p \equiv$???
N3. *A* accepts $H!(\text{not-}p) \equiv A$ hoorays that not-*p*.

(N1) follows from (E) by contraposition, and (N3) is a substitution-instance of (E). However, there can plainly be no similar analysis of (N2), since the sentence '*A* hoorays that *p*' does not admit of an intermediate form between external and internal negation. Presumably what must happen in this case is that *A* has some other sort of attitude towards *p*. A mere absence of attitude is evidently not enough, since we shall otherwise fail to distinguish (N1) from (N2), which is where Blackburn's original definition goes wrong. Acceptance of the negation of a sentence must involve something positive, something more than just a commitment not to do something else. Yet it is unsatisfactory if *all* that can be said in case (N2) is that *A* has some alternative positive attitude towards *p*, because unless we can also demonstrate a suitable internal connection between it and the original attitude, the *logical* relationship between '$H!p$' and 'not-$H!p$' will have been lost. This is just the Frege–Geach problem as applied to negated formulae. More generally, it is not enough to have a different attitude for each logical context. Unless the attitudes are connected in some systematic way, then we shall lack anything approaching a logic or a recursive semantics.

This is why it would be wrong to suppose that the error with Blackburn's original definition of '$T!p$' is no more than a minor slip, and that all we really need is for our new formula to stand for whatever attitude of thinking permissible may be most appropriate here. True, attitudes of this kind certainly exist. It is also true that the introduction of a dual operator enables one to drive external negations inwards, which is all that is needed to deal with the surface-syntax problem. Nevertheless there remains a deeper problem, namely, that we still lack the means of actually formulating the fact that '$T!$' is to be the dual of '$H!$'. It needs to be a logical law (or, at the very least, a '*quasi*-logical law' in some sense or other) that we cannot properly hold both '$H!p$' and '$T!{\sim}p$'. However, if '$H!$' and '$T!$' are merely understood to denote expressions of two different kinds of attitude (and that is all) then, as we have just observed, this fact will be lost from sight. It becomes at best a mere brute fact that the attitudes conflict with each other, with no internal complexity that could explain why.

We are not concerned with the more familiar problem of whether joint possession of these attitudes is a genuinely logical error, a merely moral one, or something in between – which is what the Blackburn–Hale dispute is primarily about, a debate which tends to dominate the issue. We are concerned, rather, with the more basic and more urgent problem of why we should suppose there to be *any* clash (of any kind). Blackburn's main point (quoted above), that both to promote p and to promote $\sim p$ is to be in a situation where not all one's goals can be realized, is totally unhelpful here. All that it explains is why we cannot accept both $H!p$ and $H!\sim p$, and the latter is of (N3) negation form. The relevant internal complexity is perfectly straightforward in such cases, but they are irrelevant to our problem. Even if we insist that the language of 'hooraying' (like moral language) has an essentially practical function, and that it should express our policies and not just our wishes, this will not account for all the structure we need if we are to be able to simulate realism. The role of the (N2) negation form remains unaccounted for.

It may be protested that logical (or '*quasi*-logical') relations do not always have to be realized in structural complexity of some kind. After all, 'red' and 'green' are exclusive, perhaps logically exclusive, even though there is no syntactic complexity that underlies this fact. The analogy is unpersuasive, however. There is an immediacy about seeing red and green which, *pace* G.E. Moore, has no ethical counterpart. Indeed, the logical relationships between the basic deontic categories, for example, between the obligatory and the permissible, are highly controversial, and so can hardly be grounded in what is self-evident (which is where red–green exclusivity seems to be grounded).[9] The objection may be pressed that this shows only that syntactic structure (which leads only to comparatively uncontroversial logical relations) is not always needed in logic after all. Moreover, there are other kinds of example which might be relevant here. For instance, basic attitude concepts such as 'believes', 'knows' and 'doubts' are certainly logically connected, but without any evident syntactic explanation for this. In general, there is more to logic than formal logic; there may be a mystery as to *how* this is so, but there need be no doubt *that* it is so. However, this does not help our problem either. We are primarily concerned with what makes '$H!p$' and '$\sim H!p$' jointly unsatisfiable, and the reason for *that* must surely be syntactic! If there are no structural connections between the concepts associated with the '$H!$' and '$T!$' operators, and a non-syntactic explanation is given of why '$H!p$' and '$T!\sim p$' are jointly unsatisfiable, then this simply argues against the original translation of '$\sim H!p$' as '$T!\sim p$'. At any rate, such a translation must become very controversial: and yet it has to be utterly uncontroversial if it is to be accepted at all, since the two formulae cannot be independently compared. The need for syntactic structure remains.

[9] See, e.g., P. McNamara, 'Making Room for Going beyond the Call', *Mind*, 105 (1996), pp. 415–50.

Another line of objection is that all we have managed to show so far is that a semantics for negation does not follow immediately from the (E) schema alone; and it may be wondered why it should have been expected to do so. After all, the operator '*A* accepts that ... ', within which '*H*!*p*' is securely embedded, is hyperintensional, and such contexts are notoriously logic-resistant, even when dealing with ordinary truth-apt sentences. The difficulty, however, is in seeing where else the semantics can come from. Acceptance contexts are built into expressivist theories at a deep level, since such theories explain the meanings of moral terms first and foremost in terms of what we are doing when we accept (or express or avow or evince) sentences that use them. Insistence on the primacy of attitude over expression, of '*A* hoorays that *p*' over 'Hooray to *p*!', and of the fundamental importance of the (E) schema, amounts to little more than an admission of this point. Acceptance-contexts do indeed make logical analysis very difficult, but this does not reflect a fault in the way in which the problem has been set out. On the contrary, such contexts are indispensable, and the Frege–Geach problem is a problem for just this reason. The apparent emergence of three negatable forms (N1), (N2) and (N3) from a sentence that admits of only two remains a mystery, and does not derive from any simple misunderstanding.

III

So where else can the structure come from? For Blackburn, it in fact emerges from the way in which the model-set semantics is imposed:

> Suppose that to a standard first-order language we add operators *H*! and *T*! subject to the condition that if *A* is a well formed formula, *H*!*A* is well formed and *T*!*A* is well formed. Suppose now we start with a set of sentences *L*, which may contain sentences with these operators among them. We begin by defining a *next approximation to the ideal*, L^*, of *L*.
>
> (i) If $H!A \in L$, then $H!A \in L^*$
>
> (ii) If $H!A \in L$, then $A \in L^*$
>
> (iii) If $T!A \in L$, then a set L^* containing *A* is to be added to the set of next approximations for *L*
>
> (iv) If L^* is a next approximation to the ideal relative to some set of sentences *L*, then, if $A \in L^*$, $A \in$ subsequent approximations to the ideal L^{**}, L^{***}, ...
>
> We can say that a set of *final ideals*, $\{L^{***} \ldots\}$ of *L*, is obtained when further use of these rules produces no new sentence not already in the members L^{***} ... of the set.
>
> The set of sentences *L* may contain disjunctions or conditionals ready to be treated as disjunctions in the deductive apparatus [*sc.* a standard tableau system]. We can say that to each branch of a disjunction there corresponds a *route* to an ideal.

We can then define: a set of sentences L is unsatisfiable iff each route to a set of final ideals S results in a set of sentences S one of whose members contains both a formula and its negation.[10]

There are many further technical details here, but they need not concern us. The crucial point is simply that, according to this system, if '$H!p$' holds at an index then 'p' holds in *all* next approximations to the ideal, and if '$T!p$' holds at an index then 'p' holds in *at least one* next approximation to the ideal. It easily follows that '$H!p$' and '$T!{\sim}p$' are jointly unsatisfiable.

This appears to solve the problem. The trouble, however, is that it is far from clear why an expressivist is entitled to use this sort of semantics in the first place. In standard deontic logic, we can define 'p is obligatory' as 'p holds in all deontically perfect worlds', and 'p is permissible' as 'p holds in at least one deontically perfect world'. So defined, the logical connection between the obligatory and the permissible is obvious, and a model-set semantics will do no more than mirror this connection. This is why the use of such a semantics is justified. Yet 'p holds in all deontically perfect worlds' and 'p holds in at least one deontically perfect world' are complex, quantified sentences, not simple expressions of attitude. Blackburn glides over this point too easily, and writes that '$H!p$ can be seen as expressing the view that p is to be a goal, to be realized in any perfect world – a world in which ${\sim}p$ is less than ideal, according to this commitment' ('Attitudes and Contents' p. 189). This is how the atomic formula '$H!p$' is first introduced. Yet we cannot just assume, as Blackburn appears to, that 'p is realized in any [i.e., every] perfect world' is nothing more than a gloss on 'p is a goal', if the latter is also to be understood as nothing more than a simple expression of conative attitude. It looks as if the structure we need is being pulled out of thin air.

Standard deontic logic works on the assumption that its formulae are ordinary truth-apt sentences, and this is evidently more than a purely arbitrary preference. The only way for the expressivist to proceed here is to insist that the sentence 'p is realized in every perfect world' is, although complex, still a non-descriptive expression of attitude, since the word 'perfect' is non-descriptive. This is indeed an inevitable response, but it can hardly help us with the analysis, for the word 'perfect' is embedded within the scope of the quantifier. If we do not yet know how to solve the Frege–Geach problem as applied to simple sentential operators, then we certainly do not know how to solve it as applied to quantifiers! Should we proceed along these lines, then at a first approximation, 'p is realized in every perfect world' would have to be analysed as '$(\forall w)(Pf!w \rightarrow p \in w)$', and '$p$ is realized in at least one perfect

[10] 'Attitudes and Contents' p. 194.

world' as '$(\exists w)(Pf!w \ \& \ p \in w)$', where the quantifiers range over worlds, construed as maximal sets of sentences, and '*Pf*!' is a perfection-attitude operator. Clearly this is to invoke far more syntactic complexity than we are currently in a position to handle. Moreover, even if further development in that direction were possible, it would make the analysis of '$H!p$' considerably more elaborate than Blackburn wishes it to be. Indeed, he regards it as an advantage that his logic is highly general, applying not just to the restricted notions of the obligatory and the permissible, but also to the more fundamental notions of the needed and the desirable. The notion of hooraying something is assumed to be pretty basic, and yet this is just what it cannot be if the logic is to work.

It may be protested that the sentence 'p is realized in every perfect world' is used by Blackburn merely to talk *about* the behaviour of '$H!p$' within the system, not to give its logical form, and that the above objection derives from a confusion between object-language and meta-language. It is true that the formula '$(\forall w)(Pf!w \rightarrow p \in w)$' is open to misinterpretation in this respect, but the central point is unaffected. What matters is simply that in Blackburn's system the logical connection between '$H!p$' and '$T!p$' derives from that between 'all next approximations to the ideal' and 'some next approximations to the ideal', and this is plainly something that cannot just be taken for granted, or merely stipulated. Indeed, it will *not* so derive unless the concept of hooraying itself has a certain kind of internal complexity whose origin has yet to be explained. Regardless of how the model-set semantics is introduced, the English sentence 'p is realized in every perfect world' has a complex syntactic structure (which the formalization makes explicit), and will therefore be problematic if 'perfect' is deemed to be non-descriptive, which it evidently has to be. It makes no difference *where* the problem occurs, whether in the object-language, the meta-language, or anywhere else.

This point can be made clearer if we consider the simplest version of expressivism, that which defines an attitude operator in terms of desire *simpliciter*:

D. A accepts $D!p \equiv A$ desires that p.

We may avoid obvious objections to this idea by not translating '$D!p$' as 'It morally ought to be the case that p', or anything too similar. Instead we shall simply define the '$D!$' operator by this schema alone, and make no further assumptions about how it should be interpreted. If Blackburn's logic were as straightforward as he claims it to be, then we could define a dual formula '$D^*!p$' equivalent to '$\sim D!\sim p$', just as '$T!p$' is definable as '$\sim H!\sim p$'. This will yield a dual attitude which we may call 'desire*', such that

D*. A accepts $D^*!p \equiv A$ desires* that p.

But is there such an attitude at all? There does not appear to be. The nearest approximation would involve something like permission, yet this is to reintroduce the more specific notion of obligatoriness. Matters would be straightforward if we were allowed to read '$D!p$' as 'It is desirable that p', for then '$D^*!p$' would simply mean 'It is not desirable that not-p'. The trouble is that to accept that it is desirable that p is to do more than just to desire that p. It is rather to accept that we *ought* to desire that p, or something along these lines. An independent deontic operator has now been pressed into service, something which is not to be found in the more primitive notion of just desiring something. If this criticism is right, and it turns out that there is no attitude dual to desire, then it shows that (D) not only fails to define the moral 'ought', it fails to define anything at all! It is simply not a valid schema, for there is no formula '$D!p$' that both satisfies it and is also negatable. If it appears otherwise, then this is because the (N1), (N2) and (N3) forms of negation have not been distinguished carefully enough. Moreover, it is clear that the situation will not be changed merely by importing a Hintikka–Blackburn system of model-set semantics (together with appeals to respect the distinction between object-language and meta-language, and so forth). This is not to deny that we *could* import such a system in a purely formalistic way. It is just that such a system will not genuinely *apply* to formulae such as '$D!p$' when they are taken to mean what they are intended to mean, namely (unqualified) expressions of desire, a fact that may not be obvious if we merely study the formalism.

The simple but easily obscured fact is that '$D!p$' will lack a negation unless there is an attitude dual to desire; and if there is no such attitude to begin with, then the introduction of a Hintikka–Blackburn type of semantics will not create one. The same point applies with schemata similar to (D) which use other intentional verbs instead of 'desires', such as 'promotes', 'hopes', 'fears', 'knows', 'doubts', and so on. In fact, none of these other verbs appears to have a dual. Nor was there ever any *a priori* reason for supposing otherwise. If this is not also the case with 'hoorays', then it has yet to be explained why. We have already noted that the practical function of hooraying will not help us, since all that is needed here is a conflict between '$H!p$' and '$H!{\sim}p$', not between '$H!p$' and '$T!{\sim}p$'. We do not need a dual operator to express the fact that we cannot aim at incompatible goals. The only obvious remaining explanation, alas, is the one that has been most obvious from the outset, namely, that 'A hoorays that p' is just shorthand for 'A accepts that [some deontic operator that] p', and that the syntax of the latter sentence should be taken at face value. If this is so, then expressivism is false.

Despite these points, it may still be felt that the difficulties have been exaggerated. The central idea of Blackburn's semantics is just that the meanings of complex formulae are to be explained by their inferential roles, given some appropriate conception of inference (or '*quasi*-inference'). We may therefore wonder why negated formulae should be singled out, given that the rules of inference for negation are, if anything, rather less complex than those for conjunction and disjunction. However, other operators are affected by these difficulties as well. The rules for conjunction are indeed very simple, and nobody finds it hard to understand a conjunction of two expressions of attitude, or a conjunction between one such expression and an ordinary descriptive sentence, if only because there is only a minimal difference between the conjunction of (any) two speech acts and their mere succession. Disjunction, however, is far less straightforward, for to obtain an adequate account we need to distinguish the following (their conjunctive analogues are, of course, all equivalent):

D1. Either A accepts $H!p$ or A accepts $H!q$
D2. A accepts $H!p \vee H!q$
D3. A accepts $H!(p \vee q)$.

(D1) and (D3) are unproblematic, but (D2), which is midway in strength between them, admits of no immediately obvious expressivist analysis. The situation is clearly similar to the case of negation. This time, however, Blackburn is more aware of the problem, and his solution invokes the standard tableau treatment for ordinary disjunction:

> Thus $p \vee H!q$ issues in a branch [of a tableau]; it is the commitment of one who is what I shall call *tied to a tree* – that is, tied to (accepting that p or endorsing q), where the parentheses show that this is not the same as (being tied to accepting p) or (being tied to endorsing q). Rather, the commitment is to accepting the one branch should the other prove untenable. The essential point is that this is a quite intelligible state to be in.[11]

But is it? The difficulty is that expressive disjunction is here defined in terms of the disjunctive syllogism, which requires a negated premise; and negation has yet to be analysed. To accept that the goal that q is untenable has to be more than just to withdraw acceptance of q as a goal. Rather it has to be a case of accepting $\sim H!q$, as opposed to not accepting (or withdrawing acceptance of) $H!q$; and until it has been explained what this means, the system's entire inferential structure (apart from its purely conjunctive part, which is trivial) is jeopardized. Disjunctions and conditionals tend to hog the limelight in this debate, but it can now be seen why negated formulae present a more fundamental source of embarrassment.

[11] 'Attitudes and Contents' p. 192.

IV

Can we perhaps amend Blackburn's account slightly? It may be insisted that all these difficulties can easily be rectified by treating the notion of hooraying something as a second-order intentional concept, for example as follows:

E′. A accepts $H!p$ (i.e., A hoorays that p) $\equiv A$ desires that we aim at p.

The complexity of the right-hand side is now manifest, and we certainly obtain three negatable forms. Moreover the (N2) form, 'A desires that we do not aim at p', is not too implausible a rendering of 'A accepts $\sim H!p$'. On this view, the dual operator '$T!$' is definable as follows:

Def. $T!$. A accepts $T!p \equiv A$ desires that we do not aim at $\sim p$.

The source of the incompatibility of '$H!p$' with '$\sim H!p$' (i.e., '$T!\sim p$') is now very clear; someone who accepts both sentences thereby desires that we both aim and do not aim at p, and this is clearly irrational. True, it is not strictly speaking *illogical*, since it is a case of inconsistent desires rather than of inconsistent beliefs, but it is nevertheless an excellent candidate for '*quasi*-inconsistency' of the sort that Blackburn requires. Such a revised proposal concedes that a semantics for negation cannot come from a formal system imposed from without, but only from a structure to be found already within the basic concept of hooraying itself. However, the proposed structure is one which does not compromise the non-cognitive status of this concept, and this is all that the expressivist needs.

This strategy may look very attractive, even obvious. The trouble is that it does not fit in with Blackburn's overall logic. The model-set semantics still requires that if '$H!p$' ('$T!p$') holds at an index, then 'p' holds at every (at least one) next approximation to the ideal, and this new strategy does nothing to explain where the quantifier structure is supposed to come from. Structure has been introduced, certainly, but of an entirely different kind. It may be protested in reply, however, that once second-order intentionality is introduced in the manner suggested, then we shall no longer need the quantifier structure. Perhaps we should discard the model-set semantics altogether, and develop a new kind of formal semantics directly. Thus the treatment for negation becomes immediate, as shown above, and analogous treatments perhaps work for the other sentential operators, for example:

Condit. A accepts $H!p \rightarrow H!q \equiv A$ desires that if we aim at p then we aim at q

Disjunc. A accepts that $H!p \vee H!q \equiv A$ desires that we either aim at p or else aim at q.

More generally, any complex formula of the form 'A accepts that S', where 'S' includes attitude operators, would be analysed by replacing 'accepts' by 'desires', and '$H!p$' by 'We aim at p'. If such a semantics were to succeed, it would provide a uniform and spectacularly simple solution to the Frege–Geach problem, one which could be generalized to all kinds of contexts.

Alas, it does not work! We need only consider mixed formulae where expressions of attitude are combined with ordinary truth-apt sentences, such as 'p & $H!q$'. The acceptance of this formula should evidently include the acceptance that p, that is to say, the belief that p; but the solution under consideration demands instead that it include the desire that p. Admittedly, the (E′) schema is rather crude as it stands, but refinements are unlikely to help. If the first intentional verb on the right-hand side (currently, 'desires') is replaced by 'believes', then the theory will cease to be expressivist; and if it is replaced by anything other than 'believes', then the above objection will apply. Any further details are irrelevant.

The only alternative is to retain (E′) but deny that all unasserted contexts should be analysed as suggested. A promising modification is that we first put formulae into disjunctive normal form, and treat the acceptances of atomic attitude-free subformulae as ordinary beliefs and the acceptances of atomic attitude-governed subformulae (and their negations) in the way sketched above. Then the acceptance of a primitive conjunction, such as the above formula 'p & $H!q$', for example, could be analysed as a compound state in the obvious way. However, there are two difficulties. First, formulae such as '$H!(H!p \rightarrow p)$' (which Blackburn spends some time examining) become intractable, so we need to restrict the analysis to first-level formulae, i.e., those where only $H!$-free formulae fall within the scope of any $H!$-operator: such remaining formulae at least *have* disjunctive normal forms. This may not sound like a major restriction, but much of Blackburn's treatment of the important notion of convergence to an ideal makes essential use of iterated operators. The second difficulty is that, although the primitive conjunctions are easily dealt with, we still need to explain how to disjoin them. As I observed earlier, Blackburn's notion of being 'tied to a tree' involves the disjunctive syllogism and hence the negation of one of the disjuncts, and I have only explained how to negate atomic formulae. Even the negation of a simple conjunction will yield difficulties, since on the present proposal it can only be understood, via de Morgan's laws, as the disjunction of its negated conjuncts – which brings us back to disjunction again. For example, the formula '$(H!p \,\&\, H!q) \vee (H!r \,\&\, H!s)$', itself not unduly complex, will be intractable.

A natural response is to insist on using conjunctive, not disjunctive, normal forms, and to present formulae instead as conjunctions of primitive disjunctions. The trouble with this approach is that mixed primitive disjunctions cannot be analysed in any way analogous to the way their conjunctive versions can be. Thus although 'A accepts $H!p \vee H!q$', for example, may still be analysed as 'A desires that we either aim at p or else that we aim at q', 'A accepts that $p \vee H!q$' lacks any obvious direct analysis of the kind our modified system seemed to promise. In fact this difficulty will also affect the earlier proposal to use disjunctive normal form.

What further alternatives are there? It is notable that the second-order commitment involved in being 'tied to a tree' applies more naturally to conditional rather than to disjunctive forms, if only because we do not have the negation problem. In effect, '$S \vee T$' is defined as '$\sim S \rightarrow T$', and commitment to a conditional is understood directly as the second-order commitment of being committed to the consequent if committed to the antecedent. Perhaps we could avoid the negation problem if, instead of using ordinary disjunctive normal form, we replaced disjunctions with conditionals and ensured that negations continue to govern only simple $H!$-formulae or their $H!$-free subformulae (negations will be driven inwards by systematically replacing '$\sim(S \rightarrow T)$' by '$S \& \sim T$', and '$\sim(S \& T)$' by '$S \rightarrow \sim T$' throughout). Such a revised normal form is ungainly, however, and left-bracketed nested conditionals are not easily grasped even in normal circumstances. Worse, and more fundamentally, we still lack a clear understanding of these second-order commitments. Specifically, we lack a convincing analysis of 'A accepts that if S then T' (itself a straightforwardly descriptive sentence) which applies to all sentences 'S' and 'T', regardless of whether they themselves are descriptive (wholly or in part). Blackburn's account suggests something along the lines of

Condit′. A accepts $S \rightarrow T \equiv A$ ties the acceptance of S to the acceptance of T

which is highly implausible when 'S' and 'T' are descriptive sentences, since the meaning of 'If S then T' will (in general) have nothing to do with acceptances (i.e., beliefs). Yet we cannot just replace 'acceptance' by 'content of the acceptance' here, even though that is what we are aiming at, since it is precisely the idea of 'tying' one content to another that we do not yet understand in those cases where the content is non-descriptive. It could be protested that further analysis is not all that essential here, and that it is always open to us to treat the conditional form as basic. However, we shall clearly lose explanatory power if we adopt that route; and it will leave Blackburn's own objectives seriously underachieved, since it is the analysis of moral conditionals (and the legitimacy of moral *modus ponens*) that was his original topic.

Perhaps an alternative solution can be found; after all, there are indefinitely many other modifications to be explored. As things stand, however, we have not succeeded in integrating the most naturally revised proposal for negation within Blackburn's overall strategy, and we are left with something that, at best, applies only to a very limited class of formulae. Although expressivism has not been refuted decisively, it looks as if the surface-syntactic properties of moral sentences will have to be taken largely at face value after all; in which case the prospects for Blackburn's *quasi*-realist project, as originally conceived, look rather bleak.[12]

Bolton Institute

[12] Earlier versions of this article were read at the 1997 Annual Conference of the British Society for Ethical Theory held at Lancaster University, at the Philosophy Seminar at the University of Sheffield and at the Stapledon Society at the University of Liverpool. I am grateful to many people for their support and encouragement, including Simon Blackburn, Peter Carruthers, Julian Dodd, Allan Gibbard, Christopher Hookway, James Lenman, Paul McNamara and Nicholas Nathan. I am also grateful to Bolton Institute for some teaching remission in the academic years 1997–9 during which this article was developed.