Attitudes and Contents

Simon Blackburn

Ethics, Vol. 98, No. 3. (Apr., 1988), pp. 501-517.

Stable URL:
http://links.jstor.org/sici?sici=0014-1704%28198804%2998%3A3%3C501%3AAAC%3E2.0.CO%3B2-M

Ethics is currently published by The University of Chicago Press.

Attitudes and Contents*

Simon Blackburn

GENERAL CONSIDERATIONS

G. F. Schueler's paper puts in a forceful way various reservations about my treatment of indirect contexts, on behalf of the position I have called "quasi-realism."[1] His opposition is, I think, as complete as could be: it is not only that my treatment has been incomplete, which I happily concede, or that its formulation has been defective, which I am prepared to believe, but also that nothing like it could possibly succeed. That at least is the proper consequence of some of his views—on logical form, and on validity, and on the nature of commitment. For example, if to show that an inference has "the logical form" or "is an instance" of modus ponens involves taking it as "the realist picture" has it, then no attempt to explain it in other terms will be compatible with its having that form. Again, if validity is ("as it is used in logic") defined in terms of the impossibility of premises being true and conclusions false, then persons reluctant to apply truth and falsity to any of the elements of an inference will have to admit that the inference is not valid, as the term is used in logic. Third, if "talk of 'commitments' is problematic for the antirealist" then antirealism will make no headway by thinking of a more general class of commitments than those with representative or realistic truth conditions. Fortunately, none of these contentions seems to me correct. Since the survival of quasi-realism even in spirit demands their rebuttal, I shall start by considering them in turn.

1. It is not too clear what it is for an argument to have the logical form of modus ponens. If that is a remark about syntactical form, then obviously having that logical form is compatible with any number of deep and different semantics for the components. To show this compare "*P*, $P \rightarrow Q$, so *Q*" with the implication taken as truth-functional, with the same seeming argument taken as some suppose the English take it: $P \rightarrow Q$ is the commitment of one who attributes a high probability to *Q* conditional upon *P*. Which is the true modus ponens? If we plump for either exclusively, we face the uncomfortable consequence that it becomes

* I am grateful to the editor for allowing me the opportunity to reply to Professor Schueler's paper.

1. G. F. Schueler, "Modus Ponens and Moral Realism," in this issue.

Ethics 98 (April 1988): 501–517

controversial whether natural English contains any inferences of the form. If we embrace both, then being of the form modus ponens is compatible with any number of deep and different *explanations* of the semantics of the components: how it comes about that we have here elements describable as true or false, or a connective properly represented by some → or other. The same point could be made with any connective: knowing even when to interpret the negation sign of a logic as meaning negation is no easy matter.[2] If quasi-realism, in the form in which I tried to develop it, is right, the "deep" semantics of a surface example of modus ponens is to be explained in a particular, and perhaps initially surprising, way. But it is modus ponens, for all that. Or, if we say it is not, then we have no effective procedure for telling when anything is.

2. Perhaps the best way to answer the restrictive view of validity is by appeal to authority. One might cite imperative logic. Or, one might cite the approach to propositional inference in terms of coherent subjective probability functions, where validity corresponds to there being no coherent function attributing a lesser probability to the conclusion than to the premises, and coherence is defined in terms of immunity to Dutch book.[3] (This is the approach which would best marry with the probabilistic view of conditionals above.) Or one could cite the view of Stig Kanger, that in interpreting a deontic logic the extension of the truth predicate to the formulae, which could equally be regarded as imperatives or expressions of attitude, is a conventional matter.[4] A further reply would draw the usual distinction between an algebraic, mathematical, pure, or uninterpreted semantics—itself sufficient to yield notions of satisfaction, validity, and completeness—versus an applied or interpreted semantics, in which the valuation clauses reflect something about the use or meaning of the connectives.[5] Formal studies are content with the first, so that truth-in-a-structure, or satisfiability defined in terms of (for instance) sets of open sets in the topology of the real line defines validity.[6] But even when we turn to the second, the question of priorities still arises. It does not go without saying that we interpret the propositional connectives by drawing on an antecedent understanding of (classical) truth and falsity. Falsity and negation go hand in hand, and it should not be obvious which is the dominant partner. The view that it is by knowing how to *use* the

2. B. J. Copeland, "What Is a Semantics for Classical Negation?" *Mind,* vol. 95 (1986).

3. Hartry Field, "Logic, Meaning and Conceptual Role," *Journal of Philosophy,* vol. 74 (1977).

4. Stig Kanger, "New Foundations for Ethical Theory," in *Deontic Logic: Introductory and Systematic Readings,* ed. R. Hilpinen (Dordrecht: Reidel, 1971), pp. 55–56.

5. M. Dummett, "The Justification of Deduction," in *Truth and Other Enigmas* (London: Duckworth, 1978), p. 293; A. Plantinga, *The Nature of Necessity* (London: Oxford University Press, 1974), pp. 126 ff.

6. H. Weyl, "The Ghost of Modality," in *Philosophical Essays in Memory of Edmund Husserl* (Cambridge, Mass.: Harvard University Press, 1940).

connectives in proofs that we come to understand them, and hence gain what understanding we have of the truth tables, is perfectly open. In Prawitz's words: "Presumably, the observational consequences that can be drawn from the assumption that a person knows the condition for the truth of a sentence can also be drawn from the assumption that he knows how to use the sentence in proofs."[7] The whole philosophy of intuitionistic interpretations of the logical constants and of those who give priority to sequent calculi and natural deduction systems opposes the simple assumption that an antecedent understanding of "representative" truth and falsity affords the only road to understanding validity. A more plausible view and one which nicely fits quasi-realism is that attributions of validity and application of the truth predicate go hand in hand: I expand on this below.

3. Schueler finds the very notion of a commitment "problematic for the anti-realist." His argument that it is one of "those terms which seem to entail a realist picture" is this: "If I am committed to, say, paying my nephew's way through school, or to the claim that a Republican succeeded Carter, then this seems something objective, forced on me by a promise I have made or other views I hold." Well it might be, if we could suitably cash the metaphor of forcing, and suitably interpret objectivity—although whether the objectivity forced by, say, promises has anything to do with realism is another matter. But equally, the commitment might not be forced by anything, like the commitment to go for a jog once a week or to improve one's golf. It does not matter, because 'commitment' I simply use as a general term to cover mental states which may be beliefs, but also those which gain expression in propositional form, but which for various reasons philosophers such as Hume, Ramsey, Wittgenstein, Stevenson, Ayer, Hare, and I have seen in terms of such things as acceptance of rules, changes in disposition, possession of attitudes, which are worth separating from beliefs. Some commitments will have nothing to do with approval: these include the change in one who accepts a rule of inference, or treats a proposition as necessary, or accords a high subjective probability to Q upon P, and so on, as well as Schueler's example of his belief that a Republican succeeded Carter. In the sphere of ethics, approval and attitude are natural terms to work with, but it would not matter if neither fitted exactly, or if better terms for the state in question existed. What is important is the theoretical issue of whether and why the state is worth distinguishing from belief, or at least from belief with representational truth conditions thought of realistically, but since I have written extensively on this elsewhere, I shall pass that over.[8]

7. D. Prawitz, "Meaning and Proofs: On the Conflict between Classical and Intuitionistic Logic," *Theoria*, vol. 43 (1977).

8. A recent statement is my "Morals and Modals," in *Fact, Science and Value*, ed. G. MacDonald and C. Wright (London: Blackwell, 1986).

FAST-TRACK AND SLOW-TRACK QUASI-REALISM

So far I have simply dissented from Schueler's reasons for general pessimism about the approach, intending to show that the ideas behind quasi-realism survive his onslaught. But he is on stronger ground in attacking the detail of my treatment. I shall turn to that after taking stock for a moment.

The problem is that of embedding of sentences which primarily express attitude, in contexts which might appear to admit only sentences which, in some contrasting way express propositions. When I say that these sentences primarily express attitude I have never intended to deny that they can be regarded as expressing beliefs or propositions. This opposition would be going beyond anything I embrace. But I do mean that the right way of theorizing about them identifies them, in the first instance, as expressing states of mind whose function is not to represent anything about the world. They express something more to do with attitudes, practices, emotions, feelings arising in contemplating some kinds of conduct, with goal seeking, with insistence upon normative constraints on conduct, and nothing to do with representing the world. In the familiar metaphor, their "direction of fit" with the world is active—to have the world conform to them, rather than descriptive or representational. I call someone who approves both of this contrast, and with this direction of theorizing, a projectivist. Projectivism may seem to be automatically opposed to the view that in saying that something is good (etc.) we give voice to a real belief about it, and it is often so introduced (as labels like 'noncognitivism' suggest). But this opposition is not automatic. Subtlety with the concept of belief, or with the concept of truth or fact, may enable the expressivist to soften this opposition. Theory may enable us to understand how a commitment with its center in the expression of subjective determinations of the mind can also function as expressing belief, or be capable of sustaining the truth predicate—properly called 'true' or 'false.' I tried to herald this development with the notion of a "propositional reflection" in the older paper "Moral Realism," and it was the point of the last pages of chapter 6, and of chapter 7, of *Spreading the Word.*[9] It means separating truth (in this application at least) from 'represents' and its allies, but nobody has ever pointed out the harm in that.

It did, however, seem to me that, before this happy result could be secured, work had to be done. It had to be shown *why* a sentence with this role could *properly* function in the ways ethical sentences do—why it sustains a fully propositional role. I now think we should distinguish a slow track to this result and a fast track. The slow track involves patiently construing each propositional context as it comes along. This is the line I took in trying to meet Geach's problem. Its advantage is that of honest toil over what might seem like theft; its disadvantage if Schueler is right

9. Simon Blackburn, *Spreading the Word* (London: Oxford University Press, 1984).

is that it does not work. But before judging that I should admit not only that it threatens to look Ptolemaic but also that it seems not to correspond to any obvious cognitive processes we go through. It is not as though construing (say) conditionals with evaluative components comes harder to us than construing them with ordinary components, and this will need explanation.

Fast-track quasi-realism would get there in better style. It would make sufficient remarks about truth to suggest that we need a comparable notion to regulate evaluative discourse (even although that is nonrepresentational) and then say that our adherence to propositional forms needs no further explanation than that. The adoption of propositional form and style meets a need because we need to share and discuss and dissent from attitudes or other stances. It involves only philosophers in error, and little more need be said. That sounds cavalier, but it was the line of, for instance, Kant and Nietzsche and probably Wittgenstein,[10] none of whom found any particular trouble in imagining the emergence of a predicate with a nondescriptive role. Nietzsche puts it roundly:

> The pathos of nobility and distance, as aforesaid, the protracted and domineering fundamental total feeling on the part of a higher ruling order in relation to a lower order, to a *below—that* is the origin of the antithesis "good" and "bad" (the lordly right of giving names extends so far that one should allow oneself to conceive the origin of language itself as an expression of power on the part of the rulers: they say "this *is* this and this," they seal every thing and event with a sound, and, as it were, take possession of it).

Perhaps our general propensity to seal things with sounds needs no detailed explanation or justification (cf. "this is nice" as a way of voicing pleasure, and immediately giving rise to compounds "if it's nice, two would be nicer," etc.). But compromises are possible: the fast track can benefit from some of the security achieved on the slow, and the slow track can make use of some of the short cuts of the fast. Or so I shall argue. Notice that, whichever track we favor, the point is to *earn* our right to propositional forms—including the use of a truth predicate. If this is done, any conventional concept of validity tags along—there is a level of analysis at which modus ponens and the rest are no different when their components are evaluative and when they are not.

10. I. Kant, *The Critique of Judgement*, trans. J. C. Meredith (London: Oxford University Press). Although Kant believes that the judgment of taste is not a cognitive judgment (p. 41) and is determined by subjective sensation, he also thinks that since we wish to demand similarity of feeling from others, we "speak of beauty as if it were a property of things" (p. 52); Nietzsche, *The Genealogy of Morals*, first essay, II; L. Wittgenstein, *Remarks on the Foundations of Mathematics* (Oxford: Blackwell, 1956), p. 163, contains a particularly clear statement of the view that statements of mathematics mislead philosophers by their descriptive form.

EMBEDDING

A parallel to the idea that a certain sentence expresses an attitude—"Hooray for the Bears"—would be the obvious truth that some others express commands—"Go to see the Bears"—and questions—"Are the Bears doing well?" Now when imperatives and questions give rise to subordinate clauses, the linguistic forms typically maintain an indication of the original mood, even if there is another syntactic change: "he told me *to* go to see the Bears, he asked me *whether* the Bears are doing well, if *I am to* go to see the Bears, I had better have some tea first." Here the right thing to say is that the subordinate clause maintains the mood of the original, but that it is not uttered with the *force* that a direct utterance of the sentence has (nothing is commanded or questioned). Nevertheless, mood is in some sense primarily an indicator of force. It is only by understanding what a question or command is that one understands the function of the interrogative or imperative mood.

There is a prima facie puzzle here. Mood is primarily an indicator of force, force is lost in subordinate clauses, but mood is not.

I do not think, however, that the puzzle is very deep, although its formal representation can be difficult. The subordinate clause in "He said that *P*" identifies which proposition he asserted; the clauses in "he told me to go to see the Bears" or "he asked whether the Bears are doing well" identify which order he gave or which question he asked. Mood indicates that a question or command is still part of the topic, even when the overall communication is not itself a question or command. Technically, I therefore agree with Michael Pendlebury that mood (or at least the presence of the indicator "to go . . . ," "whether . . .") affects the sense of such clauses: the embedding does not cancel the semantic significance of the mood indicator, which is to maintain some connection with an original command or question.[11] It can matter that a question or command is still in this sense part of the *topic*. Perhaps the nicest illustration of this is the difference between "he knew that the Bears had won" and "he knew whether the Bears had won" where the first simply gives us the content of his knowledge, but the mood in the second shows that what he is said to have known is the answer to a question—which might have been yes or no. Of course, saying that in this way a command or question is part of the topic is not implying that one was ever actually uttered—one can know the answer to questions that have never been asked.

We do not have a mood which in this way indicates that an attitude is part of the topic. The nearest approximation is in indirect reportage of wishes expressed in the optative: "would I were in Grantchester!" can perhaps be reported: he said that he would be in Grantchester, but there is at least a slight sense of strain. Normally, if I make plain to you what

11. M. Pendlebury, "Against the Power of Force: Reflections on the Meaning of Mood," *Mind* 95 (1986): 361–73.

I feel, say about the Bears, I will most probably do so using a sentence with an "expressive" predicate: "the Bears are great!" The report of what I said in indirect speech is then easy: he said that the Bears are great. According to projectivism, the item of vocabulary shows that the original utterance was expressive of attitude. In the subordinate clause, it remains to make attitude the topic just as overt mood indicators do. The person who said that the Bears are great expressed just that attitude about the Bears. Saying that this is what he did is not of course endorsing or subscribing to the view, any more than reporting a command or question involves reissuing it in propria persona.

Suppose we spoke an "emotivist" language, in which expressions of attitude wore this function on their faces. We would not have the predicative form, to keep such expressions in the indicative mood, but an ejaculatory mood, corresponding to that of "Hooray for the Bears." It would then be necessary to have a construction of subordinate clauses corresponding to words such as "that . . . ," "to . . . ," and "whether . . . ," which marks the original attitude as the topic. There seems no problem of semantic principle about this; "that!" "whether!" and so on might be introduced, so that "he said that! hooray for the Bears" tells us which attitude he expressed, "he wondered whether! hooray for the Bears" tells us which attitude he was pondering, and so on.

If natural languages have chosen not to register expressive force by a particular mood, they may have chosen to do it in other ways. And taking other cases of mood and force as our model, there might be no great difficulty about imagining it done by an expressive mood and yielding a smooth interpretation of at least some subordinate clauses. Here there is room for the compromise between fast- and slow-track quasi-realism: see how far you get in imagining an overtly expressive language developed in such ways, and diminish (even if not to zero) the gap between what it achieves and what we do with predication, and talk of truth.

A LOGIC

Geach concentrated upon the special case of the antecedent of conditionals in his original article. My suggestion involved first describing what we are up to in embedding what is primarily an expression of attitude in a context, making it intelligible that the context should have a function. Second, it involved giving sufficient semantic theory to show why we have the way we do of meeting that need. Thus in the case of Geach's original conditional, my suggestion of what we are up to involved taking up an attitude to an involvement of attitude with attitude, or attitude with belief. Such "second-order" stances seemed to me both needed in themselves and plausible candidates for the import of a conditional with evaluative elements. If we use '$\Rightarrow$' to signify the *involvement* of one mental state with another, the result was that a simple conditional "if lying is wrong, then getting your little brother to lie is wrong" came out as:

$$H!\ (\backslash B!L\backslash \Rightarrow \backslash B'GBL\backslash)\ ,$$

where the \ . . . \ notation shows that our topic is the attitude or belief whose normal expression occurs within the slashes. Involvement is not a logical notion, but neither should it seem mysterious. I tried to explain it by introducing the idea of a sensibility as a function from belief to attitude, attitude to attitude, and so on: it is what we would overtly talk about by saying things like "I really approve of *making* approval of an action depend on its consequences" or "believing that *should* increase your approval of this." Endorsing or rejecting such an involvement of commitments one with another is an important thing to do; it is therefore not surprising that we have a simple English form with which to do it.

Let us now consider modus ponens. We have

$$B!L$$
$$H!(\backslash B!L\backslash \Rightarrow \backslash B!GBL\backslash)$$
So: $B!GBL$.

Schueler and others rightly raise doubts about the kind of inconsistency in avowing the two initial attitudes and refusing endorsement of the conclusion.[12] In *Spreading the Word* I talked of a "fractured sensibility" which could not be a proper object of approval. Schueler reasonably asks why it could not be an object of approval and whether in any case this smacks more of a moral or evaluative problem than of a logical one. Yet modus ponens with these components is surely logically valid, and a proper semantics for expressions of attitude ought to explain why.

How do attitudes become things which enter into logical relationships, which *matter* in the theory of inference? It is well known that logical relationships between imperatives can be studied by thinking of joint satisfiability—seeing whether there is a consistent world in which each of a set of imperatives is obeyed. Similarly deductive relationships between norms can be studied by thinking of ideal or relatively ideal worlds in which the norms are met. If we have here the basis for a logic, it extends to attitude. For $H!p$ 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. The contrary attitude $B!p$ would rule p out of any perfect world, and corresponding to permission we can have $T!p$, which is equivalent to not hooraying $\sim p$, that is, not booing p.

Putting attitude to the fore, instead of the more usual obligations and permissions of deontic logic, promises two gains. The first is that writers on deontic logic usually interpret "Op" and "Pp" as purely propositional by making them describe what is obligated or permitted by some

12. Bob Hale, "The Compleat Projectivist," *Philosophical Quarterly* 36 (1986): 65–85, anticipates the difficulties with validity.

supposed background set of norms (the most notable exception to this generalization is Hector-Neri Castañeda).[13] But this divorces them from their ordinary expressive use, which is not to describe what some (possibly alien) system of norms yields but to insist upon or permit various things. If the apparatus of deontic logic can be taken over while this use is kept primary, so much the better. But there is another gain in taking the portmanteau term "attitude" rather than the particular, restricted notion of "obligation" and "permission." This is that the logical apparatus should apply wherever we have the idea of a goal or aim, and corresponding idea of something to be avoided, or not to be avoided. We need not be in the realm of the obligatory, or of *requirements,* but merely in that of the needed or even just the desirable. Consistency in goals is still a desideratum whose logic needs development. And in fact the deductive apparatus of deontic logic does not depend in any way on taking obligations and permissions as fully fledged deontic notions. The same structure exists if "*Op*" is interpreted as any kind of view that *p* be true ideally and "*Pp*" as any kind of toleration of *p*.

There is nothing surprising about using realization of goals or ideals as the final test for consistency. The ordinary notion of finding whether recommendations are consistent just is to imagine them carried out and see if that can be consistently done. But Schueler rightly raises a problem which might affect the extension to attitude. This is that consistency in attitude is not a particular virtue. I may wish that *p* and wish that $\sim p$ without particular shame. I may desire that *p,* and desire that *q,* but not desire that $p \,\&\, q$: I want to spend the evening at the theater, and I want to read my book, but I do not want to read my book at the theater. There is a sense in which my goals are inconsistent—they cannot all be realized—but, if this does not matter, then it is not sufficiently *like* the vice of inconsistency in belief to form the basis of a logic.

My comment on this is threefold. First of all, I think part of the objection comes from confusing desires with wishes. Inconsistent wishes may not matter because in wishing or daydreaming we are spinning fictions, and inconsistent fictions do not matter. This is because there is no connection with action. But for all that, inconsistency in real desire may matter. Incompatible and therefore unrealizable goals are bad in a way quite analogous to that in which inconsistent beliefs are. The latter cannot represent the world properly. But the former cannot represent how to behave in the world properly: they cannot mate together with beliefs, in the usual belief-desire psychological framework, to direct effective action. The man who believes that it is raining and that it is not is badly placed to act if he wants, say, to avoid getting wet. But so is the man who believes that it is raining but wants to get wet and not to get wet.

The second point to notice is that attitude and desire are capable of qualification. I may be subject to some desires, or some pressures (tiredness,

13. H-N. Castañeda, *Thinking and Doing* (Dordrecht: Reidel, 1975), esp. chap. 2.

mood) which suggest reading a book and others which suggest visiting the theater. Do I both want to read a book and want to go to the theater? It is a crude way of representing my state. Perhaps I want to read a book *inasumuch as* I am tired, want to go to the theater *inasmuch as* I like company. I feel the different pressures, but it is at least as natural to say that I don't know what I want to do as it is to say that I want to do both. I can indeed say that I would like to do both, but that takes us back to the realm of wishes (I would like them not to conflict somehow). So one way of diminishing the attraction of inconsistent desires is to remember the difference between full-scale, all-in desires, and attractions or pressures which are not yet resolved.

Even if this point were contested, a third defense is waiting. Although I have urged the advantage of thinking in terms of a catholic conception of attitude rather than of strict deontological notions, we could restrict ourselves to concepts of being for or against or neither for nor against things where consistency *does* matter. If this is a more limited range than the full spectrum of desire, this need not matter. If, for instance, it embraced only desires which one was inclined to submit to public scrutiny, or translate into practical advice, then there would be a corresponding restriction of the interpretion of notions such as 'goal' or 'ideal.' That is fine, provided the relevant attitudes satisfy the constraints when it comes to interpreting the logic in a domain such as ethics. Since ethics is at bottom a practical subject, this is to be expected.

In the usual metaphor, the direction of fit between desires and the world is opposite to that between beliefs and the world. The desire that p dictates action if it is deemed likely, but avoidable, that $\sim p$, whereas the belief that p needs abandoning if it is deemed likely that $\sim p$.[14] But since belief and desire do each have a direction of fit and a content, then each should be fitted to play a role in a logic of consistency. A person may flout the demand of consistency in complicated or demanding situations, but only at the cost of tension: his goals cannot be realized, or if he has inconsistent beliefs, the world cannot be as he represents it. It may be admirable that we sometimes get into states where we feel that tension, but this is so for belief as well as desire. It could not be admirable in general, and it could not be true in general, for these states are essentially characterized by responsibility either to the world, in the case of beliefs, or in our response to the world, in the case of attitude.

I therefore reject Schueler's contention that there is no legitimate notion of inconsistency. But it remains to be seen whether it ratifies my assault on modus ponens, or any other natural inference pattern. Meanwhile, there is another natural worry about my proposal. As it stands it yields no smooth extension to other propositional contexts. For instance, simple disjunction with an evaluative component does not yield an obvious second-order attitude. "Either Johnny has done something wrong, or

14. Michael Smith, "The Human Theory of Motivation," *Mind* 96 (1987): 36–61.

Freddy has" is not well represented as $H!(\backslash B!J\backslash$ OR $\backslash B!F\backslash)$ where "OR" introduces a kind of disjunctive relation between attitudes. Because even if the idea of a disjunctive relation between attitudes makes sense, one might know that one of them has done something wrong but quite disapprove of taking up a negative attitude to *either* of them—if neither has yet been proved guilty, for instance. The stance that $H!$ $(\backslash B!J\backslash$ OR $\backslash B!F\backslash)$ expresses seems to be that of someone who endorses only psychologies which contain at least one of the embedded attitudes (this would be the natural interpretation of disjunction), but this is not at all the same as the stance of someone who thinks that either Johnny has done something wrong or Freddy has. One could interpret disjunction by first translating it into the associated conditional and then using the account of conditionals on that. But there is something ad hoc about such procedures. They take the theory too far from anything which seems necessary for the ordinary truth-functional disjunction, and their very unnaturalness raises again the question of adequacy. Even if the notion of involvement gives a reasonable surrogate for implication, there may be no such notion naturally available in each case of potential embedding.

Suppose then we take the theory of inference as primary. If we ask what these embeddings are *for,* the immediate answer is that they mediate inference. They show us the deductive relationships between our commitments, and between our commitments and our beliefs. So rather than *replace* logical constants, as in the approach I just gave, we might try to *retain* them and to provide an interpretation of embeddings of attitude in the contexts the deductive system is to treat: in the first place, contexts provided by the truth functors.

We know, or think we know, what the negation, disjunction, and conjunction of an ordinary proposition is. It needs showing that we have any right to extend those notions to cover expressions of a different kind. Thus in the language to come $H!p$ is to be treated as a well-formed formula capable of entering the same embeddings as p. Even if this provides a language which is formally workable, it still needs showing that it provides one which is interpretable—in which $H!p$ can still be regarded as fundamentally expressive of attitude.

Consider first negation. What can $\sim H!A$ mean? Schueler might say: nothing much to do with truth or falsity, and since that reversal is the fundamental effect of negation, it cannot mean anything to apply the notion here. But I have already remarked that it might go the other way round: falsity of p is the truth of the negation of p. Ordinary negation is expressive of denial: $\sim p$ is that proposition whose expression denies p. There is a clear corresponding relation between attitudes: there exists that attitude which 'denies,' or rejects, having p as an aim or goal. If $H!p$ expresses the attitude of endorsing the goal p, $\sim H!p$ then expresses that of opposition: tolerating $\sim p$ or allowing it as consistent with an ideal world. So we can say that $T!A$ is substitutable for $\sim H!A$, and $H!A$ for $\sim T!A$. Such a conversion drives external 'negations' on attitudes inward.

So even if the original occurrence of the external negation made us uneasy, *formally* the unease is dissipated by the conversion, and *philosophically* dissipated by recognizing sufficient analogy between conflict of attitude and conflict of belief.

What of other truth functional contexts? It is an important feature of inference using propositional calculus embeddings that they can all be represented by the normal forms of conjunction and disjunction. In a tableau development each move either adds to the string (as when A, B, are appended under A & B) or divides the string (as when A makes one branch, and B another, under A v B). Now let us suppose that we are involved with an evaluative commitment, $H!A$ or $T!A$ in a propositional calculus embedding. We can see what this means if we can interpret the strings in which it issues. This reduces then to the problem of interpreting these two elements in a tree structure. But it is easy to see what a string represents if, underneath this embedding, we get $H!p$ occurring alone (e.g., under p & $H!p$ we get $H!p$). This means that the initial complex commits us to the attitude, and this is not hard to interpret. Being against clergymen and for free love commits one to being against clergymen, and the notion of consistent realization of aims or goals shows us why.

There remains the case in which a tableau under a complex branches and $H!p$ belongs to one of the branches. The interpretation is that one potential route to drawing out the consequences of the complex involves this commitment, although another may not. Thus p v $H!q$ issues in a branch; it is the commitment of one who is what I shall call *tied to a tree*. That is, tied to (*either* accepting that p, or endorsing q), where the brackets 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. Philosophically we justify the procedure by analogy with the ordinary notion of accepting a disjunction, which similarly ties one to a tree of possibilities, and formally the language admits of identical deductive procedures.

How does this relate to the original proposal for treating the conditional? Under a material conditional $A \rightarrow B$ we get the tree with $\sim A$ in one branch and B in the other. Suppose then we treat Geach's conditional "if lying yourself is wrong getting your little brother to lie is wrong" this way. Someone asserting it is tied to the tree of (either assenting to "lying yourself is not wrong" or to "getting your little brother to lie is wrong"). What was right about my original proposal is that being so tied is in this case characteristic of a particular value system or set of attitudes. Only someone with a certain view of the relation between doing things directly and doing them indirectly is apt to assent to the conditional. This represents the reason for his being tied. What was inelegant about the original proposal, I now think, was putting that directly into the content of the conditional itself. The assent to the conditional itself does not tell us why someone is tying himself to that tree—it only tells us that he is tied and

that we can use this fact in assessing the consistency of his position. I return to this below, after detailing the logic a little.

I do not want to claim finality for the semantics I shall now sketch, but it illustrates how a logic might be developed, and it shows that notions of inconsistency and satisfiability can be defined (it also bears out Kanger's remarks mentioned above). It uses Hintikka's notion of a set of "deontic alternatives."[15] In Hintikka's semantics the central notion is that of norms obtaining in a possible world, and of the deontic alternatives to that world being the possible worlds which are in accordance with those norms. Hintikka compares this notion to Kant's "Kingdom of Ends" (*Reich der Zwecke*): it represents a "mere ideal" (Kant: "freilich nur ein Ideal") which is not realized but which we nevertheless must be able to think of consistently.

In Hintikka's development we work in terms of a model system or set of model sets. A model set is a partial description of a possible world or alternative. A set of sentences including oughts and permissions will be satisfiable if it is embeddable in such a set. This means that a set of sentences L is satisfiable if and only if there is a model system S and a model set $m \in S$, such that L is a subset of m. Logical truth of A is unsatisfiability of the negation of A; B is a logical consequence of A if and only if $A \rightarrow B$ is valid, that is, $(A \,\&\, \sim B)$ is unsatisfiable.

Where we are not concerned with attitudes the notion of a model set m is defined in a standard way:

If $p \in m$, then not $\sim p \in m$;
if $p \,\&\, q \in m$, then $p \in m$ and $q \in m$;
if $p \vee q \in m$, then $p \in m$ or $q \in m$ or both;
if $(Ex)p \in m$ then $p\ (a/x) \in m$ for some individual constant a;
if $(Ax)p \in m$ and if the free singular term b occurs in the sentences of m, then $p\ (b/x) \in m$;
$p\ (a/x)$ is the result of replacing the variable x by the singular term z everywhere in p.

Henceforward I shall depart somewhat from Hintikka's terminology, in order to separate some of the main ideas more obviously. Suppose we add to a standard first-order language 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.

(Ii) If $H!A \in L$, then $H!A \in L^*$;
(Iii) If $H!A \in L$, then $A \in L^*$;
(Iiii) If $T!p \in L$, then a set L^* containing p is to be added to the set of next approximations for L.

15. J. Hintikka, "Deontic Logic and Its Philosophical Morals," in *Models for Modalities* (Dordrecht: Reidel), 1969.

(Iiv) 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. 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.

These rules need a little gloss. Obviously Iii embodies the aim that an ideal relative to a starting set of attitudes is obtained by specifying that the goals expressed are met. Rules Ii and Iii merely ensure that the attitudes specified originally remain in the subsequent realizations (if it is good that people are kind, it remains good in a world in which they are kind). The statement $T!A$ gets handled slightly differently. It sees A as compatible with perfection, but not mandatory. When tolerations are in play we have to consider both developments in which they are realized, but also developments in which they are not. One next approximation for $(T!p \ \& \ T!{\sim}p)$ should contain p, and another $\sim p$, but the fact that they are inconsistent with each other does not reflect back on the original sentence. So to assess consistency we need to think of formulae as producing a *set* of next approximations. The rule is that if $T!A$ is present in a set, there must be *a* next set in which A is present, although it is not to be in all. This means that we shall have to consider sets of next approximations to the ideal and sets of final idealizations. Intuitively, what is to matter is whether each such set is consistent.

It may be that modifications of Iiv would be desirable. What is obtained by realizing a toleration might not automatically feed through to subsequent approximations. The intuitive idea would be that something may be tolerated now, whereas were some ideal to become realized, it would no longer be tolerable: in that cast Iiv would need qualification.[16]

We want to iterate the procedure of generating a next ideal. This can be done by repeated use of these rules. If L^* is already a next approximation to the ideal and contains a sentence A, then except where A derives from realization of a toleration, it must transfer to further approximations to the ideal L^{**}. . . . This is not so in general if A belongs to an original set L (for it may be a pity that A: $A \ \& \ H!{\sim}A$ is consistent).

16. In "Moral Quasi-Realism" (forthcoming in proceedings of 1987 *Realism and Reason* Conference at St. Andrews) Hale shows that Iiv is too strong as it stands and may better be replaced with a closer version of Hintikka's original rule.

Here the idea is that once we are following out what is so in the progressive approximations to a perfect world, any realized ideal remains realized. The denizens of paradise do not move.

To get a feel for such a semantics consider the formula $H!p \rightarrow p$. This is not valid: $\{H!p, \sim p\}$ is satisfiable. The next approximation is $\{H!p, p\}^*$, and this is a final ideal and is consistent. (As the gloss of Iiv showed, p does *not* transfer through to L^*—as far as the original set goes, p may be a pity, and this is reflected in its absence from the final ideal.) Now however consider $H!$ $(H!p \rightarrow p)$. This is valid, for $T!$ $(H!p$ & $\sim p)$ is not satisfiable. By Iv $(H!p, \sim p)^*$ must be added as a next ideal. But the ideal under that is $\{H!p, p, \sim p\}^{**}$ which is inconsistent. Here is the operation of Iiii: it was already in this working out of the ideal that $\sim p$, so it stays there when we further consider the ideal obtained by realizing $H!p$ and generates inconsistency.

Transferred to these terms Hintikka's main worked example is this. Prior once took it as a "quite plain truth" of logic that

$$H!p \ \& \ (p \rightarrow H!q) \rightarrow H!q \ .$$

But $\{H!p, p \rightarrow H!q, T! \sim q\}$ is perfectly satisfiable. In tree form and using the notion $=0\Rightarrow$ to signify adding the next approximation to the ideal, we get:

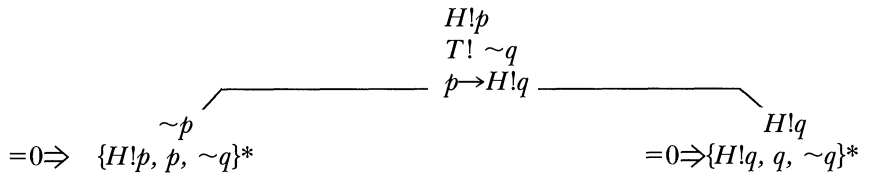

The right-hand route is bound to contain the inconsistent set, but the left yields none—reflecting the fact that if something which ought to be so is not, obligations or norms or goals consequential on its being so need not be held either. As in the first example, something here is a pity, and this is reflected by the fact that $H!$ $[H!p$ & $(p \rightarrow H!q) \rightarrow H!q]$ is indeed valid.

The possibility of valid formulae with wide scope $H!$ suggests a notion of "deontic validity" (Hintikka's term): in other words, although A may be consistent, $H!A$ need not be. In turn this gives us a needed notion: a person may be something worse than "immoral," or possessing contingently defective attitudes, but not be "inconsistent" in the sense of believing anything logically false. He may simply have ideals or goals which admit of no consistent realization.

This logic yields one reduction principle immediately: $H!H!p$ yields $H!p$. What about $T!T! \rightarrow T!p$? $\{T!T!p$ & $H!\sim p\}$ $=0\Rightarrow\{T!p$ & $H!\sim p\}$ and this too is inconsistent. This reflects the "one-dimensional" way in which realizations of goals are treated: we look through $H!$ and $T!$ to see what

happens when they are realized, and this transparency extends to iterations of them. Many complexities could be introduced at this point and in connection with iterations generally.

The semantics also generates interesting sidelights on the original proposal for treating modus ponens. Suppose we took as an example "*X* is good, if *X* is good *Y* is good, so *Y* is good." Then a treatment like my original might render it:

$$H!p;\ H!\ (H!p \rightarrow H!q)\ \text{so}\ H!q.$$

And this is indeed valid. But the satisfaction is short lived, for if we turn instead to 'Giving makes happiness; if giving makes happiness then Christmas is a good thing, so Christmas is a good thing' a parallel treatment renders it:

$$p;\ H!\ (p \rightarrow H!q)\ \text{so}\ H!q).$$

This argument is invalid. There is no way of reimporting the original p into the set of final ideals. Clearly a treatment which makes a big asymmetry between these two arguments is suspicious. However, my old proposal did not quite have this form—it involved no propositional calculus embedding of $H!p$. Now that such a form is available, obviously we shortcircuit these proposals simply to get $H!p,\ H!p \rightarrow H!q,$ so $H!q,$ and similarly for the second version. Each of these is valid.

So is Schueler right that my original proposal fails to show any inconsistency in the set containing the premises of a modus ponens inference, but a denial of the conclusion? As I mentioned, the original proposal for conditionals took seriously the idea that they create an indirect context—one where the propositions or attitudes normally expressed become, in some sense, the reference or topic of the utterance. In the present development conditionals are treated as disjunctions and broken open for example by tableau methods. Is there essential opposition here? Not necessarily. The issue is whether we can interpret endorsement of $(\backslash A \backslash \Rightarrow \backslash C \backslash)$—the original interpretation—as equivalent in strength to $(\sim A \text{ v } C)$—the place conditionals now have in the logic. Only a little leeway with 'endorsement' and 'involves' ('$\Rightarrow$') is needed; as much in fact as gives the material implication its usual right to be thought of as rendering a conditional. Say: endorsing the involvement is tying oneself to the tree. In other words tying oneself to restricting admissible alternatives to those in which $\sim A$, and those in which C. You have one or the other. And the effect of this on the theory of inference, when A or C or both are evaluative, is brought out in the model theory.

CONCLUSION

Slow-track quasi-realism will want to say that these proposals analyze or give us the logical form of the arguments we are considering. Fast-track

quasi-realism need not say this. It can say: "all this is very interesting." It shows how *little* is involved if we imagine us jumping ship—changing from an expressivist language to our normal forms. But it is unnecessary to claim that we make no jump at all. That would involve, for instance, defending the claim that negation is absolutely univocal as it occurs in $\sim H!p$ and in $\sim p$, and similarly for the other constants. But this need not be claimed. All we have is sufficient similarity of logical role to make the temptation to exploit *ordinary* propositional logic quite irresistible—and that is what we naturally do. The expressivist language serves as a model showing us why what we do is legitimate—but that may be all. This is what I meant by saying that fast-track quasi-realism can benefit from the security provided by the slow. I like this methodology. We bootstrap our way into appreciating how propositional expression, the arrival of indirect contexts, and the arrival of the truth predicate meet our needs, without in any way betraying the original, economical, metaphysical vision. At times we may have taken steps, benefiting from what are only analogies between these kinds of commitments and beliefs in order to treat the former as we treat the latter. But if so these steps are little and natural.