Reprinted from

Contemporary Philosophy
in Scandinavia

Edited by Raymond E. Olson
and Anthony M. Paul

Introduction by G. H. von Wright

THE JOHNS HOPKINS PRESS
Baltimore and London

Lars Bergström

Meaning and Morals

1

Moral philosophers are sometimes concerned with *moral* problems. (This might
not come as a surprise to anyone.) But they have also displayed a great interest
in problems about the *meaning* of those words that are typically used to express
answers to moral problems (e.g., "right," "good," "ought," and so on). Problems
of the latter kind may doubtless be of some importance, but it is not entirely
obvious how they are related to those of the former kind. In particular, it may be
wondered whether any moral conclusions can be inferred from a definition of an
ethical term. Do such definitions contain or entail moral principles? This is the
question that I propose to discuss here. It is not the only question that can be
asked about the relation between meaning and morals, but it appears to be a
rather fundamental one. Moreover, my own answer to it is different from that
which seems to be taken for granted by several prominent moral philosophers.

   Let me first mention some typical definitions of ethical terms. It might be
suggested, for example, that "good" means "pleasant" or "desired upon reflec-
tion"; that "right" means "approved by me (the speaker)" or "commanded by
God"; and that "ought to be done" means "has intrinsically better consequences
than every alternative." Some definitions of this kind may be more acceptable
than others, but this is not something that I want to discuss here. Neither need
we bother about the distinction between naturalistic and non-naturalistic defini-
tions. For the purposes of this paper we may concentrate on the following
example:

   (1)   "Right" means the same as "generally approved."

Other definitions may be treated by implication. The question that I want to

Dr. Bergström, of the University of Stockholm, is a native of Sweden.

discuss, then, is whether it is possible to derive any moral conclusion from (1). In particular, does (1) contain or entail

(2)   If an action is generally approved, then it is right,

or some similar principle? In other words, is (2) a logical consequence of (1)? This may appear to be a rather elementary question. Nevertheless, it seems to me that it needs to be answered.

It should be noticed that we do not have to pay any explicit attention to possible inferences which involve more than one premise but which are otherwise similar to the step from (1) to (2). For example, the inference

(3)   "Right" means the same as "generally approved,"
        This action is generally approved,
        _____

        Therefore, this action is right,

presents no special problem. It seems clear that (3) is valid if, and only if, (2) is a consequence of (1). This presumably is not at all controversial. But *is* (2) a logical consequence of (1)? This may seem more doubtful.

2

According to the logical intuitions of many non-philosophers it would presumably be fairly obvious that (2) follows from (1). Many people would probably say that one has to accept (2) if one accepts (1).

On the other hand, some philosophers might object to this by saying that such an inference would violate Hume's thesis that one cannot derive an *ought* from an *is*. But this argument seems to beg the question. Besides, it is not self-evident that the step from (1) to (2) is a step from *is* to *ought*. It does not seem unreasonable to suppose that someone might hold that (2) is factual or that (1) is normative. In particular, it might be argued that (1) is normative if it contains or entails (2); or that (2) is not normative if it follows from (1). So Hume's thesis does not seem to be of much help here.

Some philosophers might wish to say that (2) cannot follow logically from anything at all since it is neither true nor false. By the same token it might be maintained that (3) is a so-called practical inference and that such inferences are never logically valid. However, the thesis that something cannot be logically related to something else unless it is true or false is very far from being generally accepted. Many writers, including myself,[1] have argued that it should be rejected. Moreover, it may be doubted that (2) and the conclusion of (3) are neither true nor false. It might be held, for example, that (2) is analytic and, hence (trivially), true if (1) is true. It might even be argued that (2) follows from (1) for this very reason.

[1] See my study *Imperatives and Ethics* (Stockholm: Stockholm University, 1962), pp. 32–42.

The arguments which have been indicated above may be questioned. However, it is surely reasonable to surmise that many philosophers would in fact refuse to agree that (2) follows from (1). On the other hand, it also seems that several prominent philosophers are inclined to take the opposite view. For example, G. E. Moore writes:

> It seems sometimes to be vaguely held that when a man judges an action to be right, he is merely judging that he has a particular feeling towards it, but that yet, though he really has this feeling, the action is not necessarily really right. But obviously this is impossible. If the *whole* of what we mean to assert, when we say that an action is right, is merely that we have a particular feeling towards it, then plainly, provided· only we really have this feeling, the action *must* really be right.[2]

A few pages later he says:

> No one, I think, would be very much tempted to assert that the mere presence (or absence) of a certain feeling is invariably a sign of rightness, but for the supposition that, in some way or other, the only possible meaning of the word "right," as applied to actions, is that somebody has a certain feeling towards them.[3]

Hence, it seems that Moore would accept the view that if "right" means "approved (by X)," then a given action is right if it is in fact approved (by X). This is perhaps even more obvious from the following passage:

> Thus, if, when *I* assert an action to be right, I am merely asserting that it is generally approved in the society to which *I* belong, it follows, of course, that if it *is* generally approved by my society, my assertion is true, and the action really is right.[4]

I conclude, therefore, that it may reasonably be assumed that Moore would regard (2) as a logical consequence of (1).

Now it has been pointed out by C. L. Stevenson[5] that some of Moore's arguments are inconclusive. In particular, Stevenson shows that the definition

(4)   "X is right" has the same meaning as "I approve of X,"

does not entitle one to accept

(5)   If "X is right," said by A, is true, then X is right.[6]

But it should be noticed that Stevenson's criticism applies only to the particular inference of (5) from (4). Stevenson does not question the general assumption, which seems to be implicit in Moore's arguments, that principles like (2) and (5) may follow from definitions. As a matter of fact, he seems to hold that the definition

---

[2] G. E. Moore, *Ethics* (London: The Home University Library, 1912), pp. 92–93.
[3] *Ibid.*, p. 106.
[4] *Ibid.*, p. 108.
[5] C. L. Stevenson, "Moore's Arguments against Certain Forms of Ethical Naturalism," in *The Philosophy of G. E. Moore*, ed. P. A. Schilpp (Evanston and Chicago: Northwestern University, 1942), pp. 71–90.
[6] *Ibid.*, p. 76.

(6)   "X is right" has the same meaning as "Somebody approves of X,"

does entitle one to accept (5).[7] In other words, it seems that Stevenson would maintain that although

(7)   If someone approves of X, then X is right

is not a logical consequence of (4), it is a logical consequence of (6). Hence, he would presumably also accept the view that (2) is a logical consequence of (1).

A. C. Ewing appears to be another proponent of this view. For example, in criticizing certain naturalistic definitions of ethical terms Ewing claims:

> It is an essential feature of the moral consciousness that I realize that, if I ought to do something, I ought to do it whether others feel approval of it or not.[8]

Now he seems to hold that this moral principle is logically incompatible with a definition according to which "ought to be done" means the same as "is generally approved." For he goes on to say,

> If "what ought to be done" means "what is generally approved," general approval would have to be the only factor which ultimately counted in deciding what we ought to do, and this it certainly is not.[9]

In other words, Ewing seems to hold that the definition in question entails an unacceptable moral conclusion, namely, that something ought to be done if and only if it is generally approved, and that the definition must be rejected in order to avoid this conclusion. A few pages later he seems to argue in a similar way:

> The view that "good" means "what most people desire (or like)" is open to similar objections. Most people desire and like happiness more than great virtue, yet it does not therefore necessarily follow that the former is better.[10]

It appears that Ewing is presupposing that this would follow if "good" meant "what most people desire (or like)" or if "X is better than Y" meant "Most people desire (or like) X more than Y." In general, he seems to hold that (naturalistic) definitions of ethical terms entail (unacceptable) moral principles. Hence, he would probably say that (2) is a logical consequence of (1).

It has often been suggested that, if one accepts a naturalistic definition, then one is also committed to the view that moral or ethical problems can be solved by ordinary empirical methods. For example, R. B. Brandt writes:

> It has been suggested that "is desirable" means just "is desired by somebody." If this proposal is right, then, of course, observation can tell us what is desirable.[11]

---

[7] *Ibid.,* pp. 76–77.
[8] A. C. Ewing, *The Definition of Good* (London: Routledge & Kegan Paul, 1947), p. 62.
[9] *Ibid.*
[10] *Ibid.,* p. 65.
[11] R. B. Brandt, *Ethical Theory* (Englewood Cliffs, N.J.: Prentice-Hall, 1959), p. 152.

This view seems to involve or presuppose the assumption that naturalistic definitions entail moral principles. If Brandt's argument were expressed more explicitly it would presumably run as follows. From the definition according to which "is desirable" means the same as "is desired by somebody" we may derive the principle that something is desirable if and only if it is desired by somebody. Observation can tell us what is desired by somebody. Hence, if we accept the definition, we are committed to the view that observation can tell us what is desirable. Further support for this interpretation of Brandt's argument can be found in the following passage:

> The reason why all problems of ethics can be solved by the methods of science, if naturalism is true, is that the naturalist's definitions (like every definition) enable him to assert that some fundamental ethical statements are true by definition—statements he can use as the basic premises of his system of normative ethics. For instance, Perry's definitions enable us to say, "Any act is right if and only if it will contribute more to harmonious happiness than anything else the agent could do instead." The Ideal Observer definitions permit us to assert, as true by definition, "Anything is desirable if and only if an informed (and so on) person would want it to occur." In general, a definition will permit us to say something of the form, "Anything is *E* [ethical term] if and only if it is a *PQR.*" Then, since science presumably can tell us what will contribute most to the harmonious happiness, or what an informed person would want, and so forth, it will carry us to conclusions about what is right or good.[12]

A similar view seems to be held by W. K. Frankena. He writes:

> For example, when Perry tells us that "good" means "being an object of desire," he also tells us that we can test empirically whether X is good simply by determining whether it is desired or not.[13]

And he also says,

> If "We ought to do . . ." means "We are required by society to do . . . ," then from "Society requires that we keep promises" it follows that we ought to keep promises. It will not do to reply, as some have, that no such definitions are possible since we cannot get an Ought out of an Is, for that is to beg the question.[14]

Hence, it is surely reasonable to assume that both Brandt and Frankena would hold that (2) is a logical consequence of (1).

However, Frankena might not accept Brandt's thesis that "all problems of ethics can be solved by the methods of science, if naturalism is true." For it seems to him that naturalistic and metaphysical definitions of ethical terms

> . . . do not suffice to solve the problem of justification. If we accept a certain definition of "good," or "right," then, as we saw, we will know just how to justify judgments about what is good or right. But this means that the whole burden rests on the definition, and we may still ask how the definition is justified or why we should accept it.[15]

---

[12] *Ibid.,* p. 178.
[13] W. K. Frankena, *Ethics* (Englewood Cliffs, N.J.: Prentice-Hall, 1963), p. 81.
[14] *Ibid.,* p. 80.
[15] *Ibid.,* pp. 83–84.

But this does not mean that Frankena questions the assumption that such defini-
tions entail or contain moral principles. On the contrary, his argument seems to
rest on just this assumption. This is perhaps even more obvious when he goes
on to say,

> When Perry tries to persuade us to accept his definition of "right," he is in effect
> persuading us to accept, as a basis for action, the ethical principle that what is con-
> ducive to harmonious happiness is right. . . . He cannot establish his definition unless
> he can convince us of the principle. . . . In other words, to advocate the adoption of or
> continued adherence to a definition of an ethical or value term seems to be tantamount
> to trying to justify the corresponding moral principle. . . . Such definitions . . . turn out
> to be disguised ethical principles or value judgments which cannot themselves be
> deduced logically from the nature of things.[16]

In view of all this I am inclined to believe that many moral philosophers
would maintain that definitions of ethical terms entail moral principles. At any
rate, this view seems to be held by Moore, Stevenson, Ewing, Brandt, and
Frankena, and these writers are all very influential. As far as I can see, how-
ever, this view is unjustified. I shall try to show that (2) is not a logical conse-
quence of (1). Similar cases may be treated in the same way. Hence, if I am
right about this, one may very well accept a given definition of an ethical term
without thereby committing oneself to the acceptance of any moral principle.

My argument is fairly simple, and it does not seem to involve any contro-
versial assumptions. In particular, it will not be based upon Hume's thesis (or
upon the thesis that logical relations presuppose truth values). Neither will it be
relative to any particular interpretation of (1) and (2).

3

What does it mean to say that X is a logical consequence of Y? In order to
answer this question (which is obviously of some importance in the present
context) we must distinguish between two different cases, namely: (a) the case
where X and Y are *sentences*; and (b) the case where X and Y are *statements*.
For our purposes, the distinction between sentences and statements may be
explained as follows. A sentence is a linguistic entity, a sequence of words, which
belongs to some particular language. A statement is something that may be
expressed by a sentence. It is not a linguistic entity; it does not consist of words;
it is not a part of any language.[17] A given sentence may be used or interpreted
in many different ways. In particular, it may express different statements in
different contexts or under different interpretations. Conversely, one and the
same statement may be expressed by different sentences.

---

[16] *Ibid.*, p. 84.
[17] For the distinction between sentences and statements, see, e.g., P. F. Strawson, *Intro-
duction to Logical Theory* (London: Methuen & Co., 1952), pp. 3–4. Notice, however, that
I am using "statement" in a fairly wide sense here. In particular, I include moral and norma-
tive statements, but I am not presupposing that such statements are true or false.

Let us first consider case (b). The relation which holds between two state-
ments p and q when p is a logical consequence of q may be defined or explained
in different ways, depending upon one's choice of primitive concepts and upon
the kind of problem which one is trying to answer. In this context, we may
presuppose, as a primitive notion, the relation of inconsistency between state-
ments.[18] We may then define logical consequence in the following simple way:

(D 1) The statement p is a *logical consequence* of the statement q if,
and only if, q is inconsistent with every statement which is incon-
sistent with p.

However, this definition is not applicable in case (a). It might then be
suggested that one sentence is a logical consequence of another when the state-
ment expressed by the former is a logical consequence, in the sense of (D1), of
the statement expressed by the latter. But this will not do since a sentence may
(usually) be interpreted in more than one way. We might say that (D 1) can be
directly applied to sentences if we substitute "sentence" for "statement" through-
out in my formulation of (D 1), but we are then faced with the problem of
what it means for two sentences to be inconsistent. A reasonable answer to this
is that two sentences are inconsistent when every interpretation of them is such
that the statements expressed are inconsistent. However, unless we are prepared
to maintain that no pair of (distinct) sentences is inconsistent, we must then
distinguish between what is and what is not a (permissible) interpretation. This
is rather difficult, at least in the case of a natural language, but in the present
context we need only give a rough and partial characterization of the distinction
in question.

We may conceive of an interpretation as a function which assigns statements
to sentences in a certain way. For the sake of simplicity I shall require that the
domain of an interpretation contain every sentence (of the language). If P is a
sentence and i is an interpretation we may use the expression "$i(P)$" to denote
the statement which is expressed by P according to i. Indeed, given (D 1) and
the notion of an interpretation, we may then define logical consequence between
sentences as follows:

(D 2) The sentence P is a *logical consequence* of the sentence Q if, and
only if, $i(P)$ is a logical consequence of $i(Q)$ for every interpreta-
tion i.

But this definition would be much too narrow if *any* function from sentences to
statements were an interpretation. There are at least two further conditions which
must be satisfied. First, an interpretation must not involve any deviation from
the standard meaning of purely logical terms like "not," "and," "or," "if . . .,

---

[18] For a discussion of inconsistency between statements, see, e.g., Strawson, *Introduction
to Logical Theory*, p. 2 ff.

then . . .," and so on. Secondly, it must satisfy some consistency requirement of the following kind: it does not imply that one and the same expression has different meanings when it occurs in two different sentences. This particular consistency requirement may be somewhat too strong for natural languages,[19] but we need not go into this problem here. For our purposes it will suffice to say that a given function is an interpretation if it satisfies the conditions mentioned above, but we need not exclude the possibility that some interpretation does not satisfy the last one. (It might be held that an interpretation should also be 'reasonable' from the point of view of common usage. I have no objection to this requirement, but it is usually omitted when one is interested in the purely logical relations between sentences. It seems that (D 2) would be wider or more liberal than usual if we add this requirement. However, for the sake of argument I shall not pay any attention to unreasonable interpretations in the sequel.)

The definitions (D 1) and (D 2) are probably not controversial. As far as I can see they are quite in accordance with the definitions which can be found in modern textbooks. For example, p is a logical consequence of q according to (D 1) if, and only if, q entails p according to P. F. Strawson's definition of entailment.[20] And (D 2) is very similar, in the relevant respects, to the definition of logical consequence which is offered by P. Suppes.[21]

<div align="center">4</div>

Let us now return to our main problem. Is (2) a logical consequence of (1)? It seems that (1) and (2) may be regarded either as statements or as sentences, and I believe that it will be illuminating to distinguish between these two cases. I shall consider both.

Suppose, first, that (1) and (2) are statements. We may then use (D 1). It turns out that (2) is a logical consequence of (1) only if (1) is inconsistent with every statement that is inconsistent with (2). But it seems clear that this condition is *not* satisfied. (It might be objected that (2) is a moral statement and that moral statements are not inconsistent with any statement. If this were correct, then of course the condition would be trivially satisfied. But I shall assume that moral statements can be inconsistent with some statements. This assumption is surely very reasonable.) Consider, for example, the following statement:

(8)   Some actions which are generally approved are not right.

It seems clear that (8) is inconsistent with (2). Almost everyone would presumably agree to this. Moreover, (8) is not inconsistent with (1). In order to see this (which might be less obvious to some people) one should notice two

---

[19] This has been pointed out to me by Mr. Bengt Hansson of the University of Lund.

[20] Strawson, *Introduction to Logical Theory*, p. 20.

[21] P. Suppes, *Introduction to Logic* (Princeton, N.J.: D. Van Nostrand Company, 1957), p. 68; see also pp. 21–22 and 67.

things. First, (1) says something about the expressions "right" and "generally approved," but (8) does not say anything at all about these expressions. Conversely, (8) says something about actions, but (1) does not say anything about actions. The two statements are about wholly different subjects. Second, (1) does not say anything about the expressions which occur in (8). For, by hypothesis, (8) is a statement, and a statement does not contain or consist of any expressions at all. In particular, then, (1) does not say or imply that (8) is self-contradictory. At most, it says or implies that a certain sentence cannot be used to express (8). Neither does (8) say anything about (1). In short, there is no relation between (1) and (8) which prevents us from accepting both. We are therefore entitled to conclude that they are consistent. Hence, if (1) and (2) are statements, then, since at least one statement is inconsistent with (2) but not with (1), (2) is not a logical consequence of (1).

Let us then assume that (1) and (2) are sentences. Then, by (D 2), (2) is a logical consequence of (1) only if every interpretation *i* (of these sentences) is such that the statement expressed by (2) according to *i* is a logical consequence of the statement expressed by (1) according to *i*. But this condition is not satisfied either. For example, consider an interpretation according to which (1) expresses the statement that "right" means the same as "generally approved" and (2) expresses the statement that if an action is generally approved, then it is right. Such an interpretation does not seem unreasonable. But we have already seen that the latter statement is not a logical consequence of the former. Hence, even if (1) and (2) are sentences, (2) is not a logical consequence of (1).

It should be noticed that I do not want to deny that there may be *some* interpretation relative to which the statement expressed by (2) is a logical consequence of the statement expressed by (1). For example, if (1) expresses the statement that an action is right if and only if it is generally approved, and if (2) expresses the statement that an action is right if it is generally approved, then of course, with this particular interpretation, the statement expressed by (2) follows logically from the statement expressed by (1). But (2) does not follow from (1).

<div align="center">5</div>

Has something gone wrong here? Let us consider a possible objection. It might be argued that (2) is analytic given (1), and that (2) must therefore be accepted by anyone who accepts (1). This might then be taken to show that (2) is a logical consequence of (1).

But this argument is surely invalid. Suppose, first, that (1) and (2) are statements. In this case I would be inclined to say that (2) *cannot* be analytic, since analyticity is a property of sentences. (Note, that the usual explanations of analyticity are in terms of synonymy, or the meaning of expressions, or linguistic rules.) However, we may perhaps say that a statement can be analytic in the

sense that it is empty or that it follows logically from every statement. *If* (2) were analytic in this sense, then of course (2) would follow quite trivially from (1); but it does not follow from (1) that (2) *is* analytic. What follows from (1) is rather the following statement:

> (9)   The sentence "If an action is generally approved, then it is right" is analytic;

but (9) is not equivalent to the statement that (2) is analytic. Neither is (9) equivalent to (2). I conclude that it is not the case that (2) is analytic given (1) when (1) and (2) are statements.

Secondly, suppose that (1) and (2) are sentences. They may then be interpreted in many different ways. On some of these interpretations (1) does indeed express a statement from which it follows that (2) is analytic. See, for example, the first interpretation mentioned in section 4. But on other interpretations (1) does not express such a statement. For example, see the last interpretation mentioned in section 4. Hence, the sentence (1) does not by itself entail or entitle us to conclude that (2) is analytic.

However, it might be held that there is something odd about those interpretations according to which (1) expresses a statement from which it follows that (2) is analytic and (2) expresses a non-empty statement which does not follow from the statement expressed by (1). I have used an interpretation of this kind in order to show that (2) does not follow from (1) if (1) and (2) are sentences; it might now be objected that such interpretations are inconsistent since, roughly speaking, they imply that (2) is both analytic and synthetic. As far as I can see, however, this objection is mistaken. Interpretations of this kind need not violate the consistency requirement indicated above (section 3). In particular, they do not imply that "right" and "generally approved" have one meaning in (1) and another in (2), for these terms do not occur in (1) at all. What occurs in (1) is rather " 'right' " and " 'generally approved.' " (These latter terms may reasonably be taken to denote or refer to "right" and "generally approved," respectively, but that is irrelevant here.) Moreover, it seems clear that we may propose or accept a certain interpretation of a given sentence without accepting the statement which is expressed by the sentence according to this interpretation. In particular, we may very well assign a statement to (1) from which it follows that (2) is analytic, without thereby committing ourselves to the view that (2) is analytic; hence, such an interpretation may consistently assign a non-empty statement to (2) which does not follow from the statement expressed by (1).

As a matter of fact, it seems correct to say that (2) is analytic given (1) only if (1) is a statement and (2) is a sentence. But (2) is apparently not a logical consequence of (1) in this case either—at least not in any ordinary sense. As far as I know, logical consequence is never conceived of as a relation between statements on the one hand and sentences on the other. Neither does it seem reasonable, even in this case, to maintain that one has to accept (2) if one

accepts (1). In particular, what does it mean to 'accept' a sentence? Of course, if "to accept (2)" is now taken to mean the same as "to accept the statement that (2) is analytic," then I have no objection, but this is surely a very peculiar use of "accept."

In short, there seems to be no reason for rejecting my earlier conclusion that (2) does not follow from (1). Moreover, since my argument can easily be adapted to similar cases, we may also conclude that no definition of an ethical term entails or contains any moral principle. Some readers may find this trivial, but we have also seen that several moral philosophers appear to hold the opposite view.