

"Ought" and "is"

A. The naturalistic fallacy

What has been said up to this point about prescribing and evaluating throws some light on what is often called the problem of the relation between facts and values. This problem consists of a number of related questions. First of all, there are questions about the way evaluative and prescriptive terms function in normative assertions, in contrast with the way empirical terms function in factual assertions. Secondly, there are questions as to the validity or invalidity of arguments in which normative conclusions are drawn from factual premises. Thirdly, there are questions about the use of scientific knowledge in justifying normative assertions. In this chapter I shall consider these various questions in light of the foregoing analysis of prescribing and evaluating. At the end of the chapter I shall try to elucidate the concept of the "autonomy" of normative discourse.

When words like "good" and "bad," "right" and "wrong" are employed in the expression of value judgments, they do not refer to empirically determinable properties of things, in the way that descriptive terms refer to such properties in factual assertions. In a factual assertion we claim that something has certain properties

240

which can be discovered by empirical procedures. In a value judgment we claim that something has a certain value, but its value is not an empirically determinable property. Yet there is always a set of empirically determinable properties *contextually implied* by a value judgment. These properties are the good-making and bad-making (or right-making and wrong-making) characteristics of the evaluatum. In order to justify a value judgment we must show that the evaluatum has these properties, for these are the properties in *virtue of which* it possesses value or disvalue. If a value judgment were fully justified when these properties were found to characterize the evaluatum, then the justification of value judgments would be identical in method with the (empirical) justification of factual assertions. To say that something was "good" or "bad" would be like saying that it was green or round. Normative words would have the same function as descriptive words: to designate the empirically determinable properties of things. But this is not the case. To show that the evaluatum has certain good-making or bad-making characteristics is not sufficient for the justification of a value judgment. For a value judgment does not claim that something has these properties. It claims that something has value or disvalue in virtue of having these properties. This claim can be made *only if one has adopted a standard or rule according to which a certain property is good-making or bad-making*. A value judgment cannot be justified without appealing to the standard or rule according to which the evaluatum is being judged.

It is for this reason that no prescriptive term is ever the equivalent of a descriptive term. We use prescriptive terms ("ought," "must," etc.) to tell a person what he is to do. In prescribing an act for him to do we are not describing the act. We would describe it only if he did not understand what it was we were telling him to do. Our *description* of an act is correct when the act has the empirical properties which we assert that it has. Our *prescription* of an act is correct (i.e., justified) only if the empirical properties of the act which are good-making or right-making make it better than any alternative act open to the agent. We can only know which properties are good-making or right-making if we know which standard or rule is being used in ranking the act in comparison with its alternatives.

A standard or rule of evaluation is thus implicitly referred to whenever we utter a value judgment or prescription. This standard or rule is what transforms an empirical property into a good-making or bad-making (right-making or wrong-making) characteristic. We know whether a value predicate is properly attributable to a subject when we know the good-making and bad-making characteristics of the subject. In order to know these we must know what its empirical properties are *and* what standard or rule determines that its properties are good-making or bad-making. No knowledge of its empirical properties alone can tell us whether it is a good thing or a bad thing.

Let us take some examples. We cannot say that a child knows it is wrong to steal if *all* he knows is that he is taking something which belongs to someone else without his permission. The child must also know that there is a rule of conduct which forbids this sort of act. Again, I do not know that a painting is a bad painting merely because I know that the artist chose his colors haphazardly and used no restraint in smearing paint over the canvas. I have no right to argue that because the painting was done in this manner the result is a bad painting. I have the right to argue this only if I accept the controlled choice of color and application of paint as an appropriate standard of good painting. Given the appropriateness of this standard (which must be established by the processes of validation, vindication, and rational choice) and given the facts about how the work was painted, however, I can then legitimately claim to know that it is a bad painting. For it is the standard which renders empirical facts about how the painting was done as evidence of its badness. In other words, it is the fact that the method of painting failed to fulfill the standard that makes it a bad-making characteristic of the painting. Without a standard, neither the goodness nor the badness of the painting can be known, no matter how much empirical knowledge about the painting a person might have.

It is these considerations that explain why the naturalistic fallacy is a fallacy, or more accurately, why the various errors which Professor G. E. Moore collectively titled "the naturalistic fallacy" are errors. I shall try to show that this is so by considering three different versions of the naturalistic fallacy. For convenience we may call them the definitional error, the deductive error, and the disagreement error.

The definitional error consists in the view that every normative term (value predicate or prescriptive term) is equivalent to some empirical or descriptive term or a set of such terms. The name "naturalism" is sometimes given to this view, and different "naturalistic" theories in ethics are distinguished according to the particular empirical term (or set of terms) which is claimed to define a normative term. Thus one naturalistic theory holds that "good" means "pleasant," another that "good" means "generally approved of in this society," another that "right" means "required by law or custom," another that "desirable" means "leads to the satisfaction of human needs," and so on. There are various ways of demonstrating the error committed by all theories of this sort. I shall give what seems to me the two clearest, using the accounts of the naturalistic fallacy presented by other philosophers. Each of these accounts is derived from Professor Moore's original exposition of the fallacy in his *Principia Ethica*.

The first way is the use of the self-contradiction test. I quote from Bernard Mayo's account of Moore's proof.

If x means the same as y , then it will be self-contradictory to say that something is x but not y ; if it is not self-contradictory to say so, then x cannot mean the same as y . Now even if it is true that what is good is always pleasant and *vice versa*, yet it is certainly not self-contradictory, in the normal usage of words, to say that something good is unpleasant or something unpleasant good; therefore good and pleasant cannot mean the same, as they must if one defines the other. And the same argument holds against all other possible definitions of 'good' in terms of something else. (B. Mayo, *Ethics and the Moral Life*, p. 71.)

Why does the same argument always hold? The answer is that a normative (evaluative or prescriptive) term does not function in a sentence in the same way that a descriptive term functions in a factual assertion. The normative term does not name a property or set of properties, while the descriptive term does. The source of the definitional error of naturalism is in thinking of all adjectives as names of properties, and all declarative statements as attributions of properties to objects. This simply overlooks the radical difference in our employment of expressions when we make normative assertions and when we make descriptive assertions. When a normative

term is predicated of a subject, the subject is not being described as having certain properties. We have seen that judging the value of something is not a simple matter of attributing empirical properties to it. Prescribing an act is likewise not attributing empirical properties to the act. It is true that in judging the value of something or in prescribing something, it is judged or prescribed *because* it has certain empirical properties. But the judgment or prescription is not logically entailed by the fact that the thing has those properties. It is logically entailed by that fact *plus* the acceptance of a certain standard or rule, in the light of which we have a pro-attitude, a con-attitude, or a neutral attitude toward the thing.

A second way of revealing the definitional error in naturalism is by the use of the senseless question test. This is stated in Professor Paul Edwards' account of the naturalistic fallacy.

... Consider *any* suggested definiens of 'good.' Let *x* be the suggested definiens. Then construct questions of these two types:

- (1) Is goodness good?
- (2) Is *x* good?

If the definiens is really synonymous with 'good' then (2), no less than (1), should be a senseless or self-answering question. . . . But in fact an investigation of any definiens that has ever been or could ever be suggested shows that (2) is not a senseless or self-answering question. 'Is happiness good?' 'Is obedience to the will of God good?' 'Is aiding the struggle for survival good?'—none of these is a self-answering question. (P. Edwards, *The Logic of Moral Discourse*. Glencoe, Ill.: The Free Press, 1955, p. 209.)

Here again the explanation of why no naturalistic definiens will work for the term "good" is that reference to a standard or rule of evaluation must always be made to show that the word "good" is correctly (justifiably) ascribed to the subject of an evaluative sentence. "Good" is a *grading* or *ranking* word in this context and as such is radically different from any word or set of words which merely performs the job of describing. The only possible substitutes for "good" in an evaluative sentence would be other positive grading or ranking terms, such as "worthy," "excellent," or "desirable." And the same argument which holds for "good" also holds for these terms. No descriptive words or phrases can be made equivalent in meaning to them.

Professor Richard B. Brandt has argued that the self-contradiction test and the senseless question test only tell us when two expressions are not *overtly* synonymous. They are not adequate tests for "covert" synonymy. According to Professor Brandt two terms are overtly synonymous when

... the person, for whose usage they are overtly synonymous, thinks after the briefest reflection (if the question is put to him) that the two terms are merely different verbal devices for saying the same thing; he recognizes them intuitively as alternate, freely interchangeable expressions . . . (R. B. Brandt, *Ethical Theory—The Problems of Normative Ethics*. Englewood Cliffs, N. J.: Prentice-Hall, Inc., 1959, p. 163.)

Thus the statement "X is good but is not approved of by my society" may not appear to a person to be self-contradictory at first glance. In this case "being good" and "being approved of by my society" are not overtly synonymous. But they might still be covertly synonymous. That is, they might actually be synonymous, but this might not be obvious to the person himself. Similarly, the fact that a person thinks it is sensible to ask "Is what is approved of by my society good?" only shows that the two terms are not overtly synonymous. The senseless question test does not exclude covert synonymy. And if the two terms *were* covertly synonymous, then "being approved of by my society" would be a correct definition of "being good," and the alleged error of defining a normative expression in empirical terms would not actually be an error. Consequently the self-contradiction and senseless question tests are not sufficient for disclosing the naturalistic fallacy. The definitional error is an error only if normative expressions can never be either overtly or covertly synonymous with empirical expressions.

This argument is correct as far as it goes. But there is a simple way to amend the self-contradiction and senseless question tests so as to include covert as well as overt synonymy. That is to require that the person *reflect carefully* about whether he would be willing to say that the given statement was self-contradictory or whether the given question was senseless. If, after careful reflection about how he ordinarily uses the two expressions, about what he applies them to and what he does not apply them to, and about the way he verifies (justifies) statements in which the expressions occur as

predicates, he is willing to assert confidently that the given statement is *not* self-contradictory and the given question is *not* senseless, then the alleged definitional error is not an error, since the two expressions would be *covertly* synonymous. (Actually what had been covertly synonymous would now have become overtly synonymous.) On the other hand, all that we have to do to show that the definitional error is a genuine error is to show that no normative expression is covertly synonymous with a set of empirical expressions under these conditions. We would have to show that the more a person reflects carefully about his use of the empirical terms which are proposed as the definiens of a normative term, the less he will be inclined to say that a sentence which predicates the definiendum of something and denies the definiens (or vice versa) is self-contradictory. Similar considerations would hold for the senseless question test. From what has already been said, it is clear that this always would happen. For the more a person reflected about his use of normative terms the more he would come to understand how they function in evaluating and prescribing, and consequently how they differ from his use of empirical terms.

So far I have been concerned with the definitional version of the naturalistic fallacy. This is logically connected with the deductive version of the naturalistic fallacy, which claims that a normative conclusion can be deduced from empirical premises alone. To show that this is a fallacy, it has been argued that the relation between any set of facts and a value judgment (or a prescription) is never analytic. No contradiction is ever involved in accepting the facts and denying the value judgment (or prescription). Thus even if I accept the fact that an act will be harmful both to myself and to others, I can consistently deny the wrongness of the act. Such wrongness is not *logically entailed* by the facts about its harmful effects. I might be perverse enough to adopt a standard such that an act having harmful effects to myself and others fulfills the standard. Such a "perversion" is not an *intellectual* or *logical* error.

Sometimes the case against the deductive version of the naturalistic fallacy is stated as follows: No argument which has a value judgment or prescription as its conclusion can be valid unless there is at least one value judgment or prescription among its premises. R. M. Hare makes this point in the rule: "*No imperative conclusion*

can be validly drawn from a set of premises which does not contain at least one imperative." (Hare, *The Language of Morals*, p. 28. Hare says that this rule is the point behind Moore's refutation of naturalism. *Ibid.*, p. 30.) Thus it is claimed that we cannot draw the conclusion "You ought to do X" from such a premise as "Doing X will help to alleviate the suffering of others." The prescriptive conclusion will only follow, it is said, if we add a prescriptive (imperative) premise such as "You ought to alleviate the sufferings of others" or an evaluative premise such as "It is right to alleviate the sufferings of others." For, it is argued, unless a prescription or value judgment were included in the premises, a person could consistently accept the (factual) premise and yet deny the (normative) conclusion.

I should like to suggest that this way of putting the case against the naturalistic fallacy is too strong. For it is possible to have a valid argument whose conclusion is a value judgment or prescription and whose premises do not include any value judgment or prescription. What the premises must include in that case is a decision-statement expressing the adoption of a standard or rule. An argument with a normative conclusion may be valid, in other words, under either of two conditions. At least one of the premises must be a value judgment or a prescription, or at least one of the premises must be a decision-statement expressing the adoption of a standard or rule. The standard or rule so adopted must be applicable to whatever is evaluated or prescribed in the conclusion. In the example given in the preceding paragraph, instead of including a value judgment or prescription among the premises, either of the following decision-statements may be included: "I adopt the alleviation of human suffering as an appropriate standard for judging act X and its alternatives," or "I adopt the rule: 'One ought to alleviate human suffering' as covering act X and its alternatives." The justification of any such decision-statement consists in showing that the standard or rule being adopted is an appropriate standard or rule for judging whatever is evaluated (or prescribed) in the conclusion. This would require the validation or vindication of the standard or rule, as stated in Chapters 3 and 5.

It will be noticed that whether the premises of an argument with a normative conclusion include a value judgment, a prescription, or

a decision-statement, they must always include at least one statement whose justification requires going beyond empirical verification. We may then formulate our opposition to the deductive version of the naturalistic fallacy as follows: An argument whose conclusion is a valued judgment or a prescription is not valid if all its premises are empirically verifiable. Stated affirmatively: An argument with a normative conclusion is valid only if at least one of its premises must be justified by means of the processes of validation, the most striking difference between normative and factual assertions. The truth of normative assertions depends on human decisions; the truth of factual assertions does not. A factual assertion is true if it corresponds to the way the world is, regardless of whether we want the world to be that way. (Of course we may decide to bring about certain changes in the world. But then our factual assertions concerning those things are true only if they correspond to the changes which actually have taken place. Our decisions do not *make* the assertions true.) A normative assertion is true, on the other hand, only because we have decided to adopt a standard or rule as applicable to what we are making the assertion about. Unless we make such a decision our assertion has no truth or falsity. And the way the world is does not logically determine what decision we must make. Our adoption of a standard or rule on which the truth or falsity of our assertion depends does not itself depend on the truth or things are. We must *decide* what ought to be the case. We cannot *discover* what ought to be the case by investigating what is the case.

A normative assertion cannot properly be said to be true, however, unless the decision to adopt a certain standard or rule is itself justified. We do not make a normative assertion true merely by deciding to adopt a standard or rule. The assertion is true only when there are *good reasons* for adopting the standard or rule and for applying it to what is being prescribed or evaluated in the assertion. These good reasons ultimately rest on a human decision (though not an arbitrary one). This is the decision to commit oneself to a whole way of life, in terms of which the value system containing the standard or rule is vindicated. This ultimate decision is not arbitrary because it must be made under the conditions of rational choice, as

set forth in Chapter 6. We make a normative assertion true only when we adopt a standard or rule which is justified, and it is justified only when the way of life that involves the adoption of the standard or rule has been rationally chosen. Thus we see that the logical basis of a factual assertion and the logical basis of a normative assertion are essentially different. The first includes no element of decision, whereas the second includes an element of decision at two stages—immediately, in the decision to adopt a standard or rule of evaluation, and ultimately, in the rational choice of a way of life. Perhaps the most fundamental error of naturalism is to overlook or deny this element of decision which underlies all normative assertions.

A third version of the naturalistic fallacy I have called the disagreement error. This is the error of believing that it is possible to resolve a disagreement concerning the value of something, or concerning whether an act ought to be done, solely by appeal to the facts about the thing or act in question. No one believes that *every* normative disagreement can *actually* be settled in this way. But the philosophical naturalist believes that this is the method for resolving any such disagreement rationally. His claim is that, assuming the disputants allow their opinions to be governed by the reasons relevant to the case, they would sooner or later come to agreement if they acquired empirical knowledge about the thing in question. Sufficient empirical knowledge might never be acquired in practice, but theoretically there would be a point at which knowledge of facts alone would yield agreement on values.

The error involved here may be seen in light of what has already been said about normative assertions. Let us suppose two people disagree about the value of an object X. Let us further suppose that their disagreement stems from the fact that they apply different standards in judging X. Given one person's standard, X is good. Given the other's, X is bad. Now such a dispute cannot be resolved merely by introducing more facts about the empirical properties of X. For these facts will only reveal to one person more good-making characteristics of X and to the other more bad-making characteristics of X. Each disputant will continue to disagree with the other's value judgment of X, even when both of them agree on all the facts about X, for each appeals to a different standard. Before one can

hope to resolve their dispute by introducing more facts, it is necessary to obtain their agreement on the standard to be applied in evaluating X. Since agreement on a standard can be obtained rationally only by a method which goes beyond empirical knowledge (to validation, vindication, and rational choice), knowledge of facts is never sufficient to ensure agreement on values. It is only where there is agreement about a standard (or rule), as well as agreement about the facts, that normative agreement is ensured. Agreement about what is the case, therefore, never by itself entails agreement about what ought to be the case.

The three versions of the naturalistic fallacy discussed so far are all concerned with the meaning and justification of value judgments and prescriptions. The underlying principle which each version of the fallacy overlooks is that empirical knowledge about what is evaluated or prescribed does not by itself provide justification of a value judgment or prescription. There must always be, in addition, the justification of a decision to adopt a certain standard or rule. Speaking broadly, science can only tell us the facts about things, not their value or worth. Knowledge of the empirical world, no matter how complete, is never knowledge of good and evil. In order to know what is good and evil *in* the world, we must bring standards and rules *to* the world. But, we may go on to ask, is it not possible that standards and rules themselves may be discovered *in* the world? If this *were* possible, then philosophical naturalism would be saved. For it would mean that empirical knowledge is sufficient to ensure rational agreement on values. If we acquired enough knowledge of the world, we would finally be able to discover the true standards and rules that apply to the things we wish to evaluate or the acts we wish to prescribe. Knowledge of good and evil would be empirical knowledge and hence philosophical naturalism would be shown to be true. But this attempt to save naturalism must fail, if the account I have given of normative discourse is correct. To say that standards and rules can be discovered *in* the world is to commit the naturalistic fallacy at a higher level. The naturalistic fallacy, in other words, can be applied not only to the justification of value judgments and prescriptions but also to the justification of standards and rules. We have seen that no value judgment or prescription can be deduced from factual statements alone, since we must always

appeal to a standard or rule in justifying the judgment or prescription. It must now be shown that no standard or rule can be deduced from factual statements alone and that the attempt to do so commits the naturalistic fallacy at this second level.

The error in trying to base standards and rules on purely empirical grounds becomes clear when we recall how standards and rules are justified. Their justification does not consist in some sort of empirical confirmation, but in the threefold process of validation, vindication, and rational choice. Value judgments and prescriptions themselves can be empirically verified *if* we accept a standard or rule. (This "if" is the key to understanding why the naturalistic fallacy at the first level is a fallacy.) When we ask further whether the standard or rule itself can be empirically verified, we have already gone a step beyond the verification of a value judgment or prescription. That verification can take place only under the assumption that a certain standard or rule is applicable, and we are now questioning the grounds for this assumption. That the grounds for this assumption cannot consist in empirical statements alone should be clear from the analysis of justifying standards and rules set forth in Chapters 3, 5, and 6. A standard or rule must be validated in terms of a value system we have adopted, or else it must be vindicated as a supreme norm of the system. Adopting the system as a whole can only be vindicated in terms of a way of life. Empirical verification cannot be used at this level. A way of life must in turn be justified by a rational choice, and here again we must go beyond empirical verification. It is therefore erroneous to believe that a standard or rule can be based on empirical grounds alone. No standards or rules can be discovered in the world. We must already have adopted them in order to evaluate or prescribe things in the world, and our adoption of them is not justified by reference to the way the world is.

After considering all these ways in which the naturalistic fallacy may occur, it will perhaps be thought that scientific knowledge has little or nothing to do with our knowledge of good and evil. In order to obtain a more balanced perspective on the true role of the sciences in our evaluative and prescriptive knowledge, I shall consider in the next section those aspects of the justification of normative assertions which are dependent upon the kind of enlightenment which only the sciences can give us.

B. *Values and scientific knowledge*

What is the role of the sciences in the justification of value judgments and prescriptions? Let us begin with the first step in the four-fold process of justification. Here we are already given a standard or rule and all that we need to find out is to what extent something fulfills or fails to fulfill the standard, or whether an act is in accordance with or violates the rule. Usually the knowledge we have gained from everyday experience and the techniques of thinking we use to solve everyday problems are sufficient for this purpose. We do not require the specialized knowledge and techniques of the sciences to determine, for example, the extent to which a house we plan to buy fulfills our standards of comfort and convenience, nor do we have to depend on the sciences to decide whether a person's action is in accordance with or violates the rule of keeping one's promises.

Still, there are at least two important kinds of standard whose application often requires us to use scientific knowledge and techniques. The first kind are standards of instrumental value applied in judging the effectiveness of means to certain ends. If our end is recovery from a serious disease, for example, we must rely on the doctor's judgment of the best remedy. This in turn depends upon the findings of such sciences as chemistry and physiology. Judging the instrumental value of means to "technological" goals also requires the use of scientific knowledge and techniques. To judge the best location for building a dam which will provide electricity and also prevent floods in a certain area requires a knowledge of physics, geology, meteorology, and other sciences. An engineer's value judgment that certain materials are best for building the dam directly or indirectly makes use of the findings and methods of a number of sciences. When we are trying to determine the best means for achieving long range social ideals—such as establishing international peace, ending racial discrimination, or providing a high standard of living for all mankind—we must again rely on many different sciences, including the "human" or "behavioral" sciences of psychology, sociology, anthropology, political science, and economics.

Standards of contributive value are the second kind of standard

whose application sometimes requires the use of scientific knowledge and techniques. If one were to judge the contributive value of a certain part of a machine to the successful functioning of the whole, technical knowledge might well have to be obtained. Similarly, various "human" sciences would be required for the verification of a judgment concerning the contributive value of an individual soldier's decisions as part of an entire military operation. There are, of course, judgments of contributive value whose verification does not depend on scientific knowledge, such as judging the aesthetic contribution of sculpture to the overall beauty of a cathedral, or judging the importance of impartiality as an element in the moral life.

Whether it is obtained from science or from common sense, however, knowledge of facts is sufficient for the verification of value judgments and prescriptions. For it is an empirical question whether a given object fulfills or fails to fulfill a given standard, or whether a given act complies with or violates a given rule. It is only when the standard or rule itself is brought into question that we must go beyond empirical verification. Let us next consider the place of scientific knowledge in the *validation* of standards and rules.

In Chapter 3, I presented three methods by which the task of validation may be accomplished. What role does scientific knowledge have in these methods? It is often helpful and sometimes indispensable in the second and third methods. Method II, we recall, is used to determine whether a standard or rule includes in its range of application the class of comparison of a given value judgment. This is done by judging the effects of fulfilling a standard or acting in accordance with a rule, as compared with the effects of failing to fulfill a standard or violating a rule. In the case of a standard, it is necessary to predict the probable consequences that would follow if it were adopted by many people and they tried to fulfill it. Such prediction may require the employment of procedures in the behavioral sciences and may rely on the findings of these sciences. Similarly, in order to judge the effects of people's generally following a rule or their generally violating it, one must make predictions and the behavioral sciences again can be of help here. When we wish to justify social practices by Method II, scientific knowledge and techniques become very important. The rules which define a

social practice are judged by their instrumental value or disvalue to a given end, and extremely complex predictions must be made of the effects of an individual's or a society's engaging in the social practice. The ethical theory of restricted utilitarianism operates on this principle, and in so far as scientific knowledge and techniques are needed to make accurate predictions of the consequences of social practices, this ethical theory depends upon the sciences (though not *solely* upon the sciences) for the justification of moral rules.

The use of scientific knowledge and techniques for predicting the consequences of human acts is also a part of Method III. When this method is used for deciding whether an exception to a rule can legitimately be made, the effects of following the rule in a given set of circumstances are evaluated as being better or worse than the effects of violating the rule in those circumstances. Here the sciences may be needed to make accurate predictions of these effects.

One can now readily understand the place of science in the third step of justifying value judgments and prescriptions, namely vindication. We saw in Chapter 5 that vindicating a value system consists in showing that the consequences of a person's adopting the system would be in basic harmony with the ideals of a rationally chosen way of life and that the system itself was part of that way of life. Scientific knowledge and techniques can enable us to make accurate predictions of the consequences of a person's adopting a value system and so make it easier for us to determine its instrumental and contributive value in realizing the ideals of a way of life.

The role of the sciences in the rational choice of a whole way of life has already been explicitly considered in Chapter 6. There the necessary and sufficient conditions of a rational choice were given, and we saw that the second condition, that of enlightenment, included a complete and accurate description of the ways of life among which the choice is being made, a complete and accurate prediction of the effects of living according to them, and a complete and accurate account of the necessary means for bringing them about. In order to obtain this intellectual enlightenment, the findings of the sciences have to be relied upon. Scientific knowledge, however, is not sufficient. Philosophical knowledge is also required, as

well as "imaginative" and "practical" knowledge of the different ways of life.

The sciences function in the justification of value judgments and prescriptions, then, as a helpful and sometimes necessary instrument for successfully carrying out the different steps of verification, validation, vindication, and rational choice. But neither the whole four-step process nor any of the individual steps is *solely* dependent upon the procedures and findings of the sciences. (The first step is solely dependent upon *empirical* procedure and findings, but these need not all be scientific.) Thus it is misleading to say that value judgments and prescriptions can be "based on" the sciences. And it is certainly false to say that knowledge of what is good or bad and of what we ought or ought not to do is scientific knowledge. Yet, as I have tried to show in this section, the grounds of our value judgments and prescriptions do very often include empirical statements established by the sciences. These statements do not logically entail the value judgments or prescriptions for which they are grounds, since the process of reasoning by which value judgments and prescriptions are justified is not pure deduction. Furthermore, scientific statements alone are not sufficient for such justification. The naturalistic fallacy is not committed when the grounds of value judgments and prescriptions are scientifically established in the aforementioned ways.

C. *The autonomy of normative discourse*

In the remainder of this chapter, I wish to point out that the logic of normative discourse is independent of the logic of factual and of other kinds of discourse. When we consider the general nature of normative discourse, we see that it is "autonomous" in four respects. I shall refer to these as follows: (1) the autonomy of normative reasoning, (2) the autonomy of normative statements, (3) the autonomy of normative truth, and (4) the autonomy of normative meaning.

1. By the autonomy of normative reasoning, I mean the fact that the rules of inference (or canons of valid reasoning) that govern the justification of value judgments and prescriptions are *sui generis*.

They set up a realm of discourse which, taken as a whole, is distinct from other realms of discourse—such as those of the sciences, of mathematics, of history, and of metaphysics and theology. The rules according to which we give good reasons for (or against) a value judgment or a prescription are logically independent of the rules governing the justification of other kinds of statement. It is true that some of the ways of reasoning which are appropriate in other realms of discourse are also appropriate in normative discourse, but the *overall structure* of normative thinking is both different from and independent of other systems of thought.

We have seen how empirical knowledge can enter into the justification of normative statements and how we must sometimes rely upon such knowledge in carrying out the steps of verification, validation, vindication, and rational choice. But the total justificatory process remains unique and distinct. To engage in it is to engage in a certain kind of practice (which is both a social activity and a thought activity), defined by certain rules. These rules are the canons of valid reasoning which determine the appropriateness or inappropriateness of different ways of thinking in the context of normative discourse. When we engage in the practice by following the rules, we thereby separate ourselves as thinkers from other "engagements." We commit ourselves to normative reasoning rather than to scientific, mathematical, or some other kind of reasoning. Consequently to engage in normative reasoning is not to engage in scientific reasoning; even though from time to time we must rely on the latter in carrying out some particular step of our normative reasoning.

2. The second trait of normative discourse that constitutes a sign of its independence from factual discourse is the autonomy of normative statements, that is, the autonomy of value judgments and prescriptions themselves. By this I mean that the *premises* (as distinct from the rules of inference) of a deductive argument with a normative conclusion may not consist solely of factual statements. As we saw in discussing the deductive version of the naturalistic fallacy, no normative statement is logically entailed by a factual statement or set of factual statements.

3. The third sense in which normative discourse is autonomous concerns the truth of normative statements. Here its autonomy con-

sists in the fact that no normative statement is empirically verifiable (or falsifiable). Procedures of empirical verification may be used in justifying a value judgment or prescription in any of the four steps of justification, and all value judgments and prescriptions can be verified empirically by reference to a given standard or rule. But neither of these points allows us to say without qualification that normative statements are empirically verifiable. The first point already implies that empirical verification does not constitute the *whole* of the justification of a normative statement, and the second point explicitly states that only when a standard or rule is given can a normative statement be verified at all. We have seen that it is always appropriate to question the justifiability of accepting a standard or rule; this is not the case with empirical or factual statements.

All three of these senses of the "autonomy" of normative discourse are closely related to each other. The autonomy of the rules of reasoning that govern the justification of normative statements accounts for the second and third instances of "autonomy." Thus we know that a normative statement cannot be deduced from factual premises alone, because we know that this kind of inference is excluded by the rules of reasoning governing normative discourse. Similarly, it is in light of the entire process of justifying normative statements, as defined by those rules of reasoning, that we see that normative statements are not empirically verifiable.

4. The fourth sense in which normative discourse may be said to be autonomous concerns the meaning (or use) of normative terms rather than the truth of normative statements or the methods by which they are justified. The autonomy of normative meaning refers to the fact that the jobs for which normative language is ordinarily and correctly employed are fundamentally different from the jobs for which factual language is ordinarily and correctly employed. Our typical use of such evaluative terms as "good," "right," "desirable," "excellent," and our typical use of such prescriptive terms as "ought" and "must," are clearly distinct from our typical use of words in describing, reporting, predicting, or explaining something. The use of evaluative terms in expressing value judgments was examined in Chapter 2, and the use of "ought" sentences for the purpose of prescribing was examined in Chapter 7. We saw that

expressing a value judgment is a matter of telling someone how we grade or rank an object as a result of our evaluation of it. We may or may not be making a recommendation to the person or guiding his choice when we utter the evaluative sentence. In the case of prescriptive terms, on the other hand, we are always recommending that an act be done and so directly guiding a person's choice. In none of these evaluative or prescriptive uses of words are we asserting a matter of fact or merely conveying information about the empirical properties of things. We are playing a different "language game." Normative discourse is thus distinguishable from descriptive discourse, not only by the logical rules governing the justification of statements, but also by the semantical rules governing the proper or correct employment of the typical words which constitute such discourse.

The *proper* or *correct* employment of expressions in any realm of discourse is determined by the normal use of such expressions. When words are used in such a way that the hearer (or reader) is not misled or confused, and he has learned the "language game" in which the words normally function, then such usage is proper or correct. In the next chapter I shall consider in greater detail what a language game is and how it helps to elucidate the concept of the proper or correct use (meaning) of a term. For our present purposes it is sufficient to point out that the rules which define the language game in which words and sentences are used for evaluative and prescriptive purposes are different from those which define the language games in which we use words and sentences for descriptive or explanatory purposes. To learn how to play one game is not to learn how to play the others, since the rules are different in each case. Indeed, we do not first have to learn how to play any non-normative language game in order to learn how to play the normative language game. Or to put it another way, we do not first have to know how correctly to employ expressions for descriptive purposes in order to know how correctly to employ expressions for uttering a value judgment or prescribing an act to someone. It is in this sense that normative discourse is autonomous with regard to its meaning. These, then, are four ways in which the autonomy of normative discourse may be understood. It should be noted that I have been speaking of normative discourse as a whole. I have not been con-

cerned with the differences between moral discourse, aesthetic discourse, political discourse, religious discourse, and other "universes" of normative discourse. All of these different languages are normative in so far as they are ordinarily used in expressing value judgments and prescriptions and in reasoning for or against value judgments and prescriptions. As such, they are all autonomous in the four ways discussed above. Their autonomy is something which they have in common with one another and which differentiates all of them from the various kinds of nonnormative discourse. But according to what criteria is each of these normative languages or "universes" of discourse differentiated from the others? This is a question which will concern me in the next part of this book. I have already dealt with it briefly in Chapter 4, but incompletely so. My next task will be to take a new look at the distinction among normative languages and to fill in some gaps in my previous account.