PRENTICE-HALL PHILOSOPHY SERIES
Arthur E. Murphy, Ph.D., *Editor*

PAUL W. TAYLOR
Department of Philosophy, Brooklyn College

Normative Discourse

PRENTICE-HALL, INC. *Englewood Cliffs, N.J.*, **1961**

published in *Mind*; the editors of *The Philosophical Review* for excerpts from articles published in that journal; Mr. Peter Winch for excerpts from his book *The Idea of a Social Science and Its Relation to Philosophy*; Harvard University Press for an excerpt from *The Structure of a Moral Code* by John Ladd; Cambridge University Press for excerpts from *An Examination of the Place of Reason in Ethics* by S. E. Toulmin; The Clarendon Press, Oxford, for excerpts from *The Language of Morals* by R. M. Hare; Cornell University Press for excerpts from *The Moral Point of View* by Kurt Baier; Basil Blackwell Ltd. and the Macmillan Company for excerpts from *Philosophical Investigations* by Ludwig Wittgenstein; Penguin Books Ltd. for excerpts from *Ethics* by P. H. Nowell-Smith; The Free Press of Glencoe, Illinois for excerpts from *The Logic of Moral Discourse* by Paul Edwards; Appleton-Century-Crofts, Inc. for excerpts from *Readings in Ethical Theory*, edited by Wilfrid Sellars and John Hospers; and Prentice-Hall, Inc. for excerpts from *Ethical Theory* by Richard B. Brandt. It should be noted that in all direct quotations italics are given exactly as they appeared in the original.

PAUL W. TAYLOR

Brooklyn, New York

Contents

1 The process of evaluation

A. Value judgments and the process of evaluation

The word "evaluate" is subject to the process-product ambiguity. It may refer to the process of carrying out an evaluation, appraisal, or criticism of something, or it may refer to the product of such a process. The process of evaluation consists in trying to determine the value of something. As a product or outcome of that process, an evaluation is a settled opinion that something has a certain value. The aim of the process is to come to a decision; the process terminates when the decision is reached. During the process we are uncertain, trying to make up our minds. When our minds are finally made up, we have arrived at an evaluation (appraisal, criticism). I shall refer to this evaluation, the product of the process, as a value judgment.

A value judgment may be formed without being uttered in an evaluative sentence, but it must always be possible in principle to utter such a sentence. To have come to a decision implies that one can answer the question, "What have you decided?" We may answer this question either by saying "I have decided to do such-and-such" or by saying "I have decided that such-and-such is the case." Only the second sort of answer is pertinent here, since the evaluation process leads to a settled opinion that something is the case; it does not (directly) lead to a decision to act in a certain way. In Chapter 2, I shall examine the nature of a value judgment and its expression in an evaluative utterance. In this chapter my purpose is to make clear what it means to carry out an evaluation. What makes a thought process an evaluative process? If we answer merely that its aim is a value judgment and that it terminates in a value judgment, we have not said what the process itself consists in. We have not answered the question: How do we think when we are evaluating something? This is the question I shall now try to answer.

The process of evaluation is either a process of *grading* something or a process of *ranking* something in comparison with other things. That which is graded or ranked I shall call "the evaluatum." Every grading or ranking process may be analyzed in terms of the following five elements:

1. *A class of comparison* whose members are the evaluatum and all the things in comparison with which the evaluatum is graded or ranked.
2. *A norm or set of norms* according to which the evaluatum is graded or ranked.
3. *A pro-attitude, a con-attitude, or a neutral attitude* on the part of the evaluator.
4. *A set of good-making or bad-making characteristics* possessed by the evaluatum.
5. *A point of view* within the framework of which the grading or ranking is carried out.

The aim of the evaluation process is to arrive at a judgment in which we make the claim that the evaluatum either fulfills or fails to fulfill the norms. The evaluation process is not to be thought of as the psychological way by which we always in fact arrive at value judgments. Rather, it is the logical method which a rational person *would* follow if he were trying to come to a careful, reflective decision about the value of something. Indeed, an appeal to the process of evaluation itself may serve as part of the method by which we justify a value judgment, as we shall see in Chapter 3.

Every process of evaluation and every value judgment contextually implies that the norms being used are appropriate or valid. By

"contextually implies" I mean that anyone who understands that an evaluation is being carried out, or that a value judgment is being made, considers it legitimate and proper to question the appropriateness of the norms, and expects the evaluator to be able to give reasons showing that they are appropriate. Because the claim that the norms are appropriate is contextually implied by every value judgment, the *complete* justification of a value judgment requires going beyond a process of evaluation in which certain norms are appealed to. We must, in addition, justify the appeal to those norms. (What constitutes such a justification will be considered in Chapters 3, 5, and 6.)

A double claim is involved whenever we evaluate something. First there is the claim that the evaluatum fulfills or fails to fulfill the given norms, a claim explicitly made when we utter a value judgment. Second there is the (contextually implied) claim that it is valid or appropriate to apply the given norms to the evaluatum. Unless good reasons can be given in support of both of these claims, a value judgment cannot be justified. In order to determine what sort of reasons might be good reasons, the process of evaluation must be examined in detail.

There are two basic kinds of evaluation, depending on whether the norms appealed to are *standards* or *rules*. I shall examine each kind in turn.

B. *Evaluation according to standards*

1. GRADING AND RANKING

When the norms are standards, the process of evaluation may consist either in grading something or in ranking something as superior, inferior, or equal to something else. (We shall see that this distinction is inapplicable when the norms are rules. In that case, all evaluations are gradings.) We grade something as good or bad, desirable or undesirable, worthwhile or worthless. We rank something when we try to decide whether it is better or worse than another thing, whether it is as good or as bad as another thing, or whether it is the best or worst of a group of things. In both sorts of evaluation a point of view is presupposed. For example, we cannot evaluate a car as good or bad, or as better or worse than another car, unless we know what point of view we are to take. We must know what we are evaluating the car *as*. Is it to be evaluated as a piece of mechanical engineering? As an object of aesthetic contemplation? As a means of family transportation? As a symbol of social prestige? Each point of view we take makes a difference, since different standards (norms) are appropriate to different points of view. (What differentiates one point of view from another is a matter to be considered in Chapters 4 and 11.)

In grading a car as good or bad, the class of comparison is the class of all cars. If the evaluator is concerned only with passenger cars available within a certain price range, the class of comparison is accordingly restricted to all passenger cars in that price range. When a car is ranked as better or worse than, or equally as good or as bad as, another car, the class of comparison is restricted to the two cars being compared. In ranking a car as the best or the worst of a group, the class of comparison includes just those cars that make up the group.

Whether evaluations are gradings or rankings, the class of comparison varies independently of the point of view. Thus the class of comparison of different evaluations may remain constant while their points of view are changed, and the point of view may remain constant while the classes of comparison are changed. This is easily seen in the case of value rankings. One may rank painting *A* as superior to painting *B* from the aesthetic point of view, and rank painting *B* as superior from the economic point of view (that is, considering the paintings as commodities or investments). Here the class of comparison remains the same while the point of view changes. On the other hand, we may rank painting *A* as aesthetically better than painting *B* and then, still keeping the aesthetic point of view, rank two other paintings in comparison with each other. In this manner we may change our class of comparison without changing our point of view.

That points of view and classes of comparison are logically independent is as true of value gradings as it is of rankings, although this may not be apparent at first glance. When we grade a painting as good from the aesthetic point of view and then change our point of view and grade it as a poor investment, does not the class of comparison necessarily shift? When we judge the painting aesthetically, the class of comparison is the class of paintings; but from the economic point of view it is compared not only with other paintings but also with other kinds of investments. If the prospective buyer has $1000 to invest and judges the painting as a good (or poor) investment, he is presumably comparing it with other ways of investing $1000. These may include not only the buying of other paintings, but also the buying of many other kinds of things, such as real estate or government bonds. This shift does more than merely broaden the class of comparison (from paintings alone to paintings plus other things). The aesthetic evaluation of the painting may include in its class of comparison many famous paintings in museums and galleries all over the world, whereas the economic evaluation includes in its class of comparison only those paintings currently on the market and priced at $1000 or less. We see that a shift in point of view *may* be accompanied by a shift in class of comparison. But is the shift in class of comparison logically entailed by the shift in point of view? Clearly, it is not. It is quite possible for a person to evaluate a painting from the aesthetic point of view and limit his class of comparison to paintings currently on the market and priced at $1000 or less. Thus in the case of value gradings as well as in the case of value rankings, the class of comparison may remain constant when the point of view changes.

In order to make clear the difference between value gradings and value rankings, it is helpful to begin by considering the difference between two meanings of the word "good." Suppose we are trying to decide whether a certain president of the United States was a good president. Do we mean good as far as presidents usually go? Or do we mean good in an absolute sense, with an ideal president in mind? In the first case, our class of comparison is the thirty-five men who have actually been president. To say that someone was a good president in this sense means that he was *better than average*. It is to claim that he fulfilled certain standards to a higher degree than most of the other men who were president. "Good" is being used as a ranking word. In the second case, our class of comparison is not the class of actual presidents but the class of all possible (imaginable) presidents. To say that a certain president was good in this sense means that he fulfilled to a high degree those standards whose complete fulfillment would define an ideal president. "Good" is here used as a grading word. It is not possible to specify exactly to what degree the standards must be fulfilled for a man to be graded as a good president rather than as mediocre or bad. That depends on what standards one is appealing to (that is, what conception of an ideal president one has in mind), how clearly those standards are defined, to what extent the degrees to which they can be fulfilled are measurable, and how "distant" from reality is one's ideal. The distance between ideal and reality in any evaluation is a reflection of the amount of severity in the evaluator's general outlook on the matter and of the height of his idealism. If one's standards are severe enough or one's ideals high enough, the grading of all actual United States presidents will result in a judgment that *not one* was a good president—that they were all poor or mediocre in varying degrees.

Such a judgment would not be possible in a value ranking. One president must be ranked as the best and one as the worst, unless more than one are equally good (bad) and are equally better (worse) than all the others. At least one will be ranked as average (mediocre, fair, so-so), and the others will be judged good or bad in varying degrees. They cannot all be ranked as bad, for there must be some who are better than average and these are by definition good. When used as ranking words, "good" and "bad" mean *comparatively* good and *comparatively* bad. They can be defined in terms of "better" and "worse," "good" meaning "better than average" and "bad" meaning "worse than average." In order to decide whether one thing is better or worse than another, some standard or set of standards must be appealed to. In order to decide what is average, a certain class of comparison must be given. The average is then definable as what is worse than about half the members of the class of comparison and what is better than about half, when all are judged according to the given standards.

It is possible for the standards of a value ranking to be identical with those of a value grading. A ranking of the presidents of the United States from best to worst may be made according to the very same standards by which they are graded as all mediocre or bad. When that occurs, it is not a contradiction to say that a certain president was both a good one and a bad one *as judged by the same standards*. To say he was a good president is to rank him as being better than most of the presidents, and to grade him as a bad president is to claim he was too far from the ideal to be worthy of a pro-attitude on our part, and so to be worthy of being praised as a good president.

The distinction between value rankings and value gradings is independent of the point of view involved. We may rank as well as grade United States presidents from a moral point of view, a political point of view, an intellectual point of view (were they intelligent, well read, etc.?), a marital point of view (were they good husbands and fathers?), and so on. For each point of view different standards of evaluation would be used, but this would not determine whether the judgment was to be a ranking or a grading. To take another example, we may judge a painting from an aesthetic, an economic, a moral, or some other point of view. But whatever the point of view, when we judge the painting to be good we must ask ourselves: Is the painting good as compared with the paintings I have seen recently and which are still fresh in my mind, or with all the paintings I have ever seen, or with all the paintings in this exhibition, or with all the paintings in the world? Or is it good as compared with what a painting of this type *should* be? If our judgment is of the first kind it is a ranking of the painting (in one of several classes of comparison). If it is of the second kind it is a grading of the painting.

2. THE LOGIC OF THE EVALUATION PROCESS

I shall now set forth the process of evaluation as a process of reasoning. Whether an evaluation consists in ranking something in relation to something else or in grading something in relation to an ideal, it may be analyzed into a series of logical steps. This analysis is not meant to be a psychological account of what happens in our heads when we evaluate something. I am not trying to describe or explain our thought processes as they actually occur. I am trying instead to make explicit the pattern or structure of the logic of our thinking. The steps are as follows:

1. *Adoption of a standard or set of standards* for evaluating (ranking or grading) the evaluatum.
2. *Operational clarification of the standards.* This consists in a set of statements to the effect that if an object O has characteristics C,

it fulfills a certain standard S to a certain degree D; that if O has characteristics C', it fulfills S to a greater (lesser) degree D'; etc.

3. *Specifications of the class of comparison.*

4. *Determining the good-making and bad-making characteristics of the evaluatum.* (In the case of value rankings this would ideally be done for every member of the class of comparison.)

5. *Deducing,* from 2 and 4, *the degree to which the evaluatum on the whole fulfills or fails to fulfill the standards.* (In the case of value rankings this would ideally be done for every member of the class of comparison. The members would then be ranked in order according to the varying degrees to which they fulfilled or failed to fulfill the standards.)

Having completed these five steps, a person would be able to give a direct and unequivocal answer to the question: What is the value of this? His answer would express a value judgment concerning the evaluatum. Thus a value judgment is a decision or verdict, resulting from the process of evaluation, which concludes the evaluation both in a logical and in a temporal sense. It is the logical conclusion of the five steps, and when it is reached the process terminates. To reach such a decision is the very point of carrying out the evaluation.

The characteristics C, C', and so on, in step 2 are what I have called good-making and bad-making characteristics. These characteristics belong to an evaluatum (or any member of a class of comparison); it is by virtue of these characteristics that something is good or bad to some degree or other, as judged by the standards of evaluation. The evaluatum itself may have characteristics of both kinds. There may be respects in which it is desirable (its virtues or merits) and respects in which it is undesirable (its deficiencies or demerits). These may balance each other off, so that *on the whole* the object is neither good nor bad, but more or less indifferent. When it has more virtues than deficiencies, or when its virtues outweigh its deficiencies (in a sense to be explained), it is on the whole good. Under the opposite conditions it is on the whole bad.

Good-making and bad-making characteristics are relative to the standards being used. A characteristic which is good-making according to one standard may be bad-making according to another. For example, if we judge a car according to standards of ample size

and seating capacity, the fact that it has a wide wheelbase and can seat six passengers are good-making characteristics. When judged by the standard of convenience in parking, these same characteristics will be bad-making. Since different points of view entail the adoption of different standards, good-making and bad-making characteristics will vary as different points of view are taken in an evaluation. The very feature that makes an advertisement a good one from the economic point of view may make it a bad one from the aesthetic or moral point of view.

Even within the same point of view, different standards may be adopted, with the result that disagreements will arise about whether certain characteristics of an evaluatum are good-making or bad-making. Two literary critics may disagree about the worth of a new novel because they disagree about what is to count in its favor. A feature which one critic believes to be a mark of literary excellence will be considered by the other an aesthetic failure. Both are taking the aesthetic point of view, but they are appealing to different value systems. (The distinction between a point of view and a value system will be dealt with at considerable length in Chapters 4, 5, and 6. Its importance has been overlooked too often in philosophical studies of values.)

In step 2 of the process of evaluation, mention was made of the degree to which an object fulfills or fails to fulfill a standard. What does "degree of fulfillment" mean in this context? The best way to answer this is to consider, first, what we mean when we grade things as excellent, very good, quite good, fair, rather poor, very poor and so on; and second, what we mean when we rank one thing as much better than, slightly better than, or about as good as another thing. In grading football players, for example, a coach might judge that one man is an excellent passer, a fair runner, a good blocker, and a poor kicker. Such judgments may or may not be numerically measurable in terms of some kind of units of goodness or badness. As a measurement of a player's running ability, the coach might use the formula: total yards gained minus total yards lost when the player carried the ball, divided by the number of plays. The coach might then set up a scale of running ability with the following categories—excellent: ten or more yards gained; good: 5-10 yards gained; fair: 0-5 yards gained; poor: 0-5 yards lost; ready-for-the-

JV's: more than five yards lost. On the other hand, the coach might not use any measurement for grading a player's blocking ability. He might simply rely on his *intuitive* judgment, based on careful observations of the player's performance as a blocker in scrimmages and games. These intuitive judgments would be expressed in such phrases as "an outstanding blocker," "a very good blocker," or "a fair blocker," without further refinement in quantitative terms.

Similar considerations hold for value rankings. Suppose we are trying to decide which of two equally priced houses we should buy. With respect to some of the standards we use (such as the distance of the houses from our place of work, or the size of the bedrooms in each house), a quantitative judgment may be made (in terms of distance in miles or volume in cubic feet of space). With respect to other standards (such as pleasantness of appearance and convenience of the arrangement of rooms), we might intuitively judge that one house fulfills them to a greater degree than the other, without being able to say to *how much* greater a degree. This does not mean that the concept of "degree of fulfillment" is inapplicable here. It only means that degree of fulfillment is not measurable. It still makes perfectly good sense to use such ranking terms as "much better than," "slightly better than," and "about as good as" in this context.

Let us now consider what is involved in the claim that an evaluatum may simultaneously have both good-making and bad-making characteristics. This claim may be restated in a less technical way: an object may be good in some respects and bad in others. This is possible because different standards are applied to different features of the object. But what determines whether the object is good or bad *on the whole?* The answer is somewhat different for gradings and for rankings. I shall consider gradings first.

There are two factors that determine whether an object is to be graded as good or bad on the whole: the degree to which the object fulfills (or fails to fulfill) all the standards applied to it, and the relative precedence of those standards. We have seen what degree of fulfillment means. The relative precedence of standards must now be clarified, at least in a preliminary way. (Full consideration of this will be taken up in Chapter 3, in connection with the validation of standards.) When we evaluate an object by appealing to more than one standard, the standards may be organized in a system of relative precedence. To say that one standard S takes precedence over another standard S′ is to say that, when we are making an over-all judgment of its value, greater *weight* is given to the fact that an object fulfills (or fails to fulfill) S than to the fact that it fulfills (or fails to fulfill) S′. (Why this is so will be explained in Chapter 3.) In other words, if we assume that the object fulfills both S and S′ to the same degree, its fulfillment of S adds more to its over-all goodness than its fulfillment of S′. Greater weight is given to a good-making or bad-making characteristic that is determined by S than to one that is determined by S′. An example will help to clarify this. Suppose we are evaluating a vacation we have just spent abroad. We say that there were good and bad things about it. We might include among the good things the nice people we met, the fine works of art we saw, and the pleasant hotels where we stayed. Among the bad things we might count poor roads for driving, bad weather, and our trying to see too much in too short a time. Suppose we then say that, on the whole, the vacation was a very good one. I suggest that the basis for our saying this is that the good aspects of the vacation outweigh the bad. How is it that one aspect can outweigh another? The answer lies in the fact that the standards fulfilled by one aspect are higher than (i.e., take precedence over) the standards fulfilled by the other. In the case at hand, the good-making characteristics of the vacation fulfilled higher standards than those which the bad-making characteristics failed to fulfill, so the former carried more weight than the latter in the final judgment. If the bad-making characteristics had failed to fulfill higher standards than those fulfilled by the good-making ones, the vacation would have been judged, on the whole, to be bad.

The assumption made here is that the vacation fulfills or fails to fulfill each of the standards to the same degree. But if the vacation were to fulfill one standard to a greater degree than another, then the former good-making characteristic would carry more weight than the latter. This determinant of the "weight" of a good-making (or bad-making) characteristic varies independently of the other determinant of "weight," namely the relative precedence of the standards involved. If an object fulfills a lower standard to a greater degree than it fulfills a higher standard there is a counterbalancing of weights. If it fulfills a higher standard to a greater degree than a lower standard there is an increment of weight on one side. No precise measurement can be made of such changes in weight, but we can have an intuitive grasp of them. A vacation which fulfilled to a high degree the relatively low standard of good roads for driving, but did not fulfill to any great extent the relatively high standard of seeing great works of art, would not be as good on the whole as a vacation which fulfilled the second standard to a greater degree than the first. Even though no units of weight can be assigned, it is clear that the over-all judgment is affected by both factors—degree of fulfillment and relative precedence of standards fulfilled.

3. PRO-ATTITUDES AND CON-ATTITUDES

How these two factors operate in the formation of a judgment of something as good or bad on the whole, may be explained by reference to the role played by our pro-attitudes and con-attitudes in the evaluation process. Given a standard which we believe appropriate to apply in evaluating a certain object, we have a pro-attitude toward the object insofar as it fulfills the standard, a con-attitude toward it insofar as it fails to fulfill the standard. Now our pro-attitude or con-attitude *increases in strength* according to the degree to which the object fulfills or fails to fulfill the standard, and according to the relative precedence of the standard. The greater the degree of fulfillment (or failure of fulfillment) and the higher the precedence, the stronger the pro-attitude (or con-attitude). In evaluating our vacation we have a pro-attitude toward each of its good-making features and a con-attitude toward each of its bad-making features. We like, approve of, or favor the vacation in virtue of the good things about it and dislike, disapprove of, or disfavor it in virtue of the bad. If our final judgment is to be positive rather than neutral or negative, the good-making features must outweigh the bad-making, not only according to the degrees to which they fulfill the given standards, but also according to the relative precedence of the standards themselves. We use the word "good" about something when we have a pro-attitude toward it, the word "bad" when we have a con-attitude toward it, and such words as "mediocre," "fair," "all right," "unobjectionable," "so-so," "not bad," and "tolerable" when we have a more or less neutral attitude toward it. If our final judgment of the vacation is that *on the whole* it was a good one, we have a pro-attitude toward it on the whole because the total strength of our pro-attitudes is greater than the total strength of our con-attitudes. If the relative strengths of the two groups of attitudes were about equal, our final judgment would be neutral. What determines the relative strength of our pro-attitudes and con-attitudes? Insofar as we are rational (i.e., insofar as we have *justifiable* pro-attitudes and con-attitudes), the relative strength of these attitudes is determined by the relative precedence of the standards we use and the degree to which the evaluatum fulfills them.

In the case of value rankings, there are three factors that determine whether we have a pro-attitude, a con-attitude, or a neutral attitude toward the evaluatum. The first factor has already been found in value gradings—our pro-attitudes and con-attitudes are directed toward the good-making and bad-making characteristics, respectively, of the evaluatum. Here the relative strength of attitudes, based on degree of fulfillment and relative precedence of standards, plays the same role in determining our over-all attitude toward the evaluatum. But this is not sufficient for our ranking the evaluatum as better or worse than something else, or as good or bad (i.e., as better or worse than average). We must also take into account the good-making and bad-making characteristics of the other member or members of the class of comparison, in relation to which the evaluatum is being ranked. If we evaluate the object by itself, which is simply the process of grading, we may conclude our evaluation with a positive judgment and hence have a pro-attitude toward the object on the whole, only to find that, when compared to most of the other objects in the class of comparison, it is inferior. What appears by itself to be good may turn out to be worse than something else, or even the worst of a group of things. Similarly, what appears by itself to be bad may be better than something else or the best of a group of things. To decide whether the evaluatum is better or worse than something else, or is the best or the worst in the class of comparison, two factors must be taken into account: whether the evaluatum is good on the whole or bad on the whole when *graded* according to the given standards, and whether its over-all goodness or badness outweighs the over-all goodness or badness of each member of the class of comparison. If its over-all goodness outweighs that of every other good member of the class of comparison, the final ranking judgment is that it is the best of the group. If its badness outweighs that of all the over-all bad members, it is finally ranked as the worst of the group. If its over-all goodness outweighs the over-all goodness of the average members of the class of comparison, or if the average members are bad over-all, it is ranked as good (i.e., as better than average). It is also ranked as good if it is bad over-all but its badness is outweighed by the over-all badness of the average members. Though bad, it is less bad than the average and therefore to be ranked as good.

A third factor becomes important when we want to know whether the final ranking judgment expresses a positive, negative, or neutral attitude. We must know the evaluator's over-all attitude toward the whole class of comparison. It is not enough to say that a positive judgment results when the evaluatum compares favorably with (i.e., has more or stronger good-making characteristics and fewer or weaker bad-making characteristics than) most members of the class of comparison, that a negative judgment results when most members of the class of comparison compare favorably with the evaluatum, or that a neutral judgment results when the evaluatum compares favorably with about half of the class of comparison and compares unfavorably with about half. We must also take into account how the evaluator judges the class of comparison as a whole. In the example of evaluating presidents of the United States, I mentioned the possibility that a person might have such severe standards or such a high ideal of the presidency that all the actual presidents would be graded as mediocre or bad. But even under these conditions some presidents would be better than others, some would be good in the sense of better than average, and one would be the best. In ranking a president as good or as the best, however, the person would not necessarily be expressing a pro-attitude toward that president. He might not admire or approve of him at all, although his attitude would perhaps be softened by the recognition of the inferiority of most of the other presidents.

It is only when "good" is used in a grading judgment that we can

be assured it expresses a pro-attitude on the part of the speaker. If it is used in a ranking judgment, it may or may not express a pro-attitude, depending on the speaker's evaluation of the entire class of things with reference to which the evaluatum is being ranked. (We shall see later that there are other uses of "good" which are not essentially connected with a pro-attitude on the part of the speaker. I discuss these in Chapter 2.)

The fact that a value grading or ranking can result in a neutral judgment (the object is neither good nor bad, but fair or mediocre) raises the question of how such a neutral judgment differs from an ordinary statement of fact. Theories of value that define "good" and "bad" in terms of interests, attitudes, desires, and other subjective states or dispositions all emphasize the positive-negative dimension of such states or dispositions. This positive-negative dimension appears to set values off from facts, for when we say that something is good or bad our pro- and con-attitudes are involved in a way they are not when we say that something is round or heavy or green. I have been trying to show how pro- and con-attitudes do have an important place in our evaluations, yet it seems to me that we may have neutral attitudes toward things when their good-making and bad-making characteristics are of approximately equal weight. An object may be in some respects good (and insofar as it is we take a pro-attitude toward it) and in some respects bad (and insofar as it is we take a con-attitude toward it), but *on the whole* neither good nor bad. Our pro- and con-attitudes toward different features of the object cancel each other out, as it were, and leave us with a neutral attitude. We then judge the object as fair, so-so, all right, unobjectionable, and so on. This is altogether different from making a statement of fact about the object.

Neutral value judgments do not indicate that no pro- or con-attitudes have been involved in an evaluation. Indeed, such attitudes must be involved for them to become neutralized. It is the fact that pro- and con-attitudes are always part of an evaluative process that distinguishes the *grading* and *ranking* of things according to *standards* from merely *comparing* or *classifying* them according to *criteria.* We may compare cities with respect to their population, cars with respect to their speed and power, or bombs with respect to their destructive force, without evaluating them as good, bad, or indifferent. To arrange cars in an order of increasing speed and power is not necessarily to arrange them in an order of merit or desirability. A comparison is evaluative only when things are graded or ranked, that is, classified according to their *worth* or *value*. And their worth or value is determined by the degree to which they fulfill standards that one has adopted. When standards are applied in judging an object, a person has a pro-attitude or con-attitude as a consequence of the fact that the object fulfills or fails to fulfill the standards. Thus grading and ranking are very different from merely comparing or classifying.

Suppose that we always ranked apples as excellent, good, fair, or poor according to certain good-making and bad-making characteristics (defined in terms of certain standards). Now suppose that we trained a child to separate correctly all the apples he was given into the four categories *without* letting him have pro-attitudes toward excellent and good apples, a neutral attitude toward fair apples, and a con-attitude toward poor apples. The child would not be able to rank or grade the apples at all. He would know the standards only as criteria of classification. What for us are the good-making and bad-making characteristics of apples would be, for the child, simply different features of apples that tell him whether to put the apples in the box marked "excellent," the box marked "good," the box marked "fair," or the box marked "poor." The child would not understand the meaning of these words, even though he could correctly apply them. If he were to point to an apple and say (correctly) "This is a good one," he would mean "This belongs in the box marked 'good.'" He would not mean what we mean by saying "This is a good one," because his pro-attitudes would not be involved. He would not, in other words, be expressing a value judgment.

The same holds true in the case of a neutral judgment. When the child says "This is a fair apple" he is not expressing a neutral value judgment. For him the apple does not have both good-making and bad-making characteristics that cancel one another out, leaving a neutral attitude. The child has no attitude at all toward the apple. Having a *neutral* attitude is far different from having *no* attitude. In order that something be the object of a neutral attitude, it must be thought to have good-making and bad-making characteristics. What counts as a good-making or bad-making characteristic is not merely that in virtue of which an object satisfies or fails to satisfy certain tests or criteria, but that in virtue of which it fulfills or fails to fulfill certain standards *and consequently is the object of pro-attitudes or con-attitudes.*

One might want to say that evaluating something is nothing but classifying it as calling for a pro- or con-attitude simply because it belongs to one category rather than another. This analysis would be inadequate. It could not account for the logical structure of the steps which make up the evaluative process, nor could it explain the justification of value judgments as a rational process. Our attitudes are not just accidently attached to our standards. They are involved when we *adopt* our standards in the first place. It is more accurate to say that they are entailed by our standards than that they simply accompany them. This does not mean that standards and attitudes can be put in propositional form so that one can be deduced from the other. The relation between attitudes and standards will be explored in later chapters of this book (especially Chapters 3, 4, 5, and 6). For the present, some enlightenment concerning their relation will be gained by examining one of the ways in which standards may be classified.

4. INTRINSIC AND EXTRINSIC VALUE

The standards appealed to in a process of evaluation may be classified as standards of intrinsic value or standards of extrinsic value. Standards of extrinsic value may in turn be subdivided into standards of inherent value, standards of instrumental value, and standards of contributive value. These distinctions were first formulated by professor C. I. Lewis in *An Analysis of Knowledge and Valuation.* I am making a departure from Professor Lewis's position by including standards of intrinsic value among standards of evaluation. It is one of the fundamental principles of Professor Lewis's theory that expressions of intrinsic value are not value judgments. Following R. B. Perry and John Dewey, Professor Lewis holds that intrinsic value results from the act of *valuing*, which is radically different from the process of *evaluating*. I intend to show that judgments of intrinsic value are as much the outcome of a process of evaluation as are judgments of extrinsic value. Professor

Lewis's concept of valuing is identical with what I would call having a pro-attitude. The act of "valuing" an object is any act governed solely by a disposition of liking or favoring the object, of preferring it to other objects, of approving of it, of prizing it, or of having some other pro-attitude toward it. In my view, a statement that something has intrinsic value is not the direct expression of a pro-attitude (or an act of valuing). It is a value judgment, a claim that something fulfills to a certain degree a certain kind of standard —in this case, a standard of intrinsic value. A person has a pro-attitude toward the thing *in consequence of* its being found, by a process of evaluation, to have intrinsic value. The existence of a pro-attitude directed toward it does not *constitute* its value.

I do not mean to deny that the verb "to value" in its ordinary usage means to prize something, to hold something dear or precious. We ordinarily say that something is valuable *to* a person when he values it in this way. To value a book, a friendship, one's health, or a letter from a loved one is to have a certain kind of pro-attitude toward the valued thing. It is to be disposed to care for the thing, to protect it from loss or harm, and to take pleasure in possessing it or in using it. When we come to value something in this way (and when accordingly it comes to have value for us), we may or may not have good reasons for valuing it. We may value something, for instance, simply out of a blind emotional or sentimental attachment to it. Or we may value it because it satisfies a neurotic desire on our part. On the other hand, when there are good reasons for our valuing something we are *justified* in valuing it. I suggest that only one sort of reason may serve as a good reason in this case: the fact that we can justify an *evaluation* (a grading or ranking) of the valued thing as a good thing. We may justify valuing our health because it is a good means to many desirable ends. Or we may justify our valuing it simply on the ground that we enjoy being healthy. (In the latter instance, we shall see, we are evaluating health as having intrinsic value.) But however we happen to evaluate our health, we must go through a process of evaluation if we are to have good reasons for valuing it. Why is this so? Valuing something is having a pro-attitude toward it, and we are justified in having a pro-attitude toward something only when it can be shown

to be *worthy* of being liked, approved of, favored, or valued. If we like or value something that can be shown to be undesirable or bad, we are not justified in liking or valuing it. In the same manner, we can justify our having a con-attitude toward something only when we can show, on the basis of an evaluation, that the thing is undesirable or bad. *Evaluating, then, constitutes the rational ground for valuing.*

Here one relation between standards and attitudes becomes clear. Attitudes can be justified only on grounds of evaluations, and evaluations require the adoption of standards. In this respect attitudes are like decisions, actions, and choices. All of these can be justified only on the basis of evaluations. (This point will be considered in greater detail in connection with the justification of value judgments.)

To say that a thing is valued by someone is not to say that it *has value* (i.e., that it is good). Valuing is having a certain kind of pro-attitude toward something and cannot be identified with the claim that the thing has intrinsic value (or any other kind of value), for such a claim is a value judgment and must be established by a process of evaluation. It is by means of this process alone that we can determine whether something is or is not intrinsically good, that is, whether it does or does not have intrinsic value. The standards appealed to in this process are standards of intrinsic value and function in the same way as do any other standards.

For this view of intrinsic value I am partly indebted to Professor C. A. Baylis who, in his essay "Grading, Values, and Choice" (*Mind*, LXVII, 268; 1958, pp. 485-501), modifies the theory of Professor Lewis to some extent in the way I have indicated. According to Professor Baylis, the existence of a pro-attitude is "initial evidence" for the presence of intrinsic value. Under certain conditions we can judge with a high degree of probability that the object of our pro-attitude has intrinsic value.

When we judge things of certain kinds, e.g. pleasant experiences, to be intrinsically good, the best initial evidence for this that we could have, I submit, is that we find ourselves *prizing* things of that kind, i.e. liking, approving, desiring, preferring, and commending them. . . . Such evidence gives us an initial probability that what we thus prize is intrinsically good. We can increase this probability by making repeated examinations of things of the same kind under circumstances which vary just enough to guard against the kinds of cognitive error which might occur. (*Ibid.*, pp. 494-495.)

I am in agreement with Professor Baylis when he states that having a pro-attitude toward something (prizing it, liking it, etc.) is *not* a defining characteristic of intrinsic value.

> . . . The relation between the statements 'X is intrinsically good' and 'X has the observed identifying property of being prized for its own sake by Y' is synthetic rather than analytic. Neither entails the other. If the former is true there is some antecedent probability that the latter is true. If the latter is true and reasonable precautions have been taken in guarding the prizing against various possible types of errors, then the former is confirmed to some degree. (*Ibid.*, p. 497.)

Professor Baylis seems to assume, however, that the statement "X is intrinsically good" does not express a value judgment in the way that "X is extrinsically good" expresses a value judgment. I disagree with him here. It seems to me that each statement is the outcome of a process of evaluation. If I am right in this (and I shall try to defend the point below), intrinsic value and extrinsic value do not differ in their relation to pro-attitudes. In both cases there must be a pro-attitude for an evaluation to take place, but the evaluation itself is a matter of *judging* things *according to a standard*, in the one case a standard of intrinsic value and in the other a standard of extrinsic value. In dealing with intrinsic value and the three subclasses of extrinsic value, I am concerned with a fundamental way of classifying *standards* of evaluation. This classification is "fundamental" only in the sense that it is helpful in understanding what it means to evaluate and as we shall see, in understanding what it means to prescribe. There are, of course, many other ways of classifying standards of evaluation. (I shall consider eight such ways in Chapter 12.)

In order to make clear the basis of the classification at hand, it is best to begin by noting the various sorts of things that can be evaluated. I have up to now used the rather vacuous words "thing" and "object" to cover anything which could possibly serve as an evaluatum or as a member of a class of comparison. This usage was not intended to give the impression that only physical objects can

serve in this capacity. Evaluata may be physical objects ("That is a beautiful building"); but they may also be persons ("Titian was a greater painter than Raphael"), acts ("The assassination of Lincoln was a misfortune for the whole nation"), events ("It was a brilliant discovery"), situations ("It is desirable to have a doctor on duty in this plant"), states of affairs ("International peace is a worthy ideal for all nations"), rules ("The no smoking rule on subways is a good rule"), policies ("Economic assistance to underdeveloped countries is a proper use of federal funds"), practices ("Trial by jury is an admirable part of any legal system"), organizations ("That labor union is a benefit to the entire industry"), physical conditions or states ("Good health is a priceless possession"), the states of mind and dispositions of others ("There is a nasty streak of envy in him"), and the immediately felt or perceived qualities of our own experience ("The dry wine tasted just right with the dinner"). In other words, anything toward which it is possible to take a pro-attitude or con-attitude can be an evaluatum or a member of a class of comparison.

The distinction between standards of intrinsic value and standards of extrinsic value can now be made. (I am following the distinctions as they were drawn by Professor Baylis.) Standards of intrinsic value must be, first, standards used only in evaluations of the immediately felt or perceived qualities of our own experience, and must be, second, standards of nonderivative value. If an evaluatum is not a felt or perceived quality of experience, it cannot have intrinsic value. Objects, acts, situations, and persons can have extrinsic value only. A felt or perceived quality of experience may have either intrinsic or extrinsic value. However, we cannot *identify* evaluations of intrinsic value with evaluations of qualities of experience. We must bring in the second criterion: that the quality be evaluated for its nonderivative value only. The essential difference between intrinsic and extrinsic value is the difference between the derivatively good and the nonderivatively good. An object, quality, act, or person has derivative (extrinsic) value when its value is derived from the value of something else. Without the other thing's having (intrinsic or extrinsic) value it would have no value of its own. Thus automobiles have extrinsic value only, for unless people thought it was desirable to get from one place to another more quickly than animals or bicycles could convey them, automobiles would have no value. Intrinsic value, on the other hand, is nonderivative. Nothing else need have value for a thing to have intrinsic value; the thing is good in itself, apart from its leading to or contributing to the value of other things. Thus the experience of listening to music may be valued intrinsically. One has a pro-attitude toward the experience not because one has pro-attitudes toward other things which the experience makes possible, but simply because the qualities of the experience are what they are. But the orchestra, the conductor, the existence of air waves, and the possession of sense organs for hearing all have extrinsic value derived from the intrinsic value of listening to music. Without the latter, no one would evaluate positively the other things which make it possible (unless they made possible some *other* experiences having intrinsic value to someone).

The experience of listening to music may have extrinsic as well as intrinsic value. One might listen to music in order to relax, enjoying both listening and the sensation of being relaxed which results from listening. To take another example, one might enjoy playing tennis but also do it for the exercise. Thus what has intrinsic value may also have extrinsic value. ". . . To say of anything that it is intrinsically good is not to deny that it is extrinsically good. It is rather to affirm that it possesses a value over and above any extrinsic value it may have." (C. A. Baylis, *op. cit.*, p. 491.) It is possible for something to have intrinsic value and yet have extrinsic disvalue. Indeed, four possible combinations of positive and negative value may occur: a quality of experience may be both intrinsically and extrinsically valuable, both intrinsically and extrinsically disvaluable, intrinsically valuable but extrinsically disvaluable, and intrinsically disvaluable but extrinsically valuable. An example of the first would be listening to music for enjoyment and relaxation; an example of the second would be suffering pain when one is ill and incurring medical expenses at the same time; an example of the third would be a child's enjoying too much ice cream so that he becomes ill; an example of the fourth would be having to work at an unpleasant job in order to have the money to spend later.

There is *normally* only one standard of intrinsic value, namely

satisfaction in or enjoyment of the quality being experienced. The corresponding standard of intrinsic disvalue is felt dissatisfaction or unpleasantness. One quality of experience is ranked as intrinsically better than another quality of experience when the first is more pleasant on the whole (or less unpleasant on the whole) than the second. In the process of comparison, however, each quality of experience must be judged as pleasant (or unpleasant) on the whole, quite aside from its relation to the known value of other things or qualities. To say that a quality is "pleasant on the whole" means that, if we experience both pleasantness and unpleasantness in connection with it, the pleasantness outweighs the unpleasantness and as a consequence the quality has a generally positive hedonic tone for us. "Unpleasant on the whole" means the opposite, where the over-all hedonic tone is negative. When a *grading* of intrinsic value is made, the evaluator has a pro-attitude toward the evaluatum (a quality of his own experience) insofar as it fulfills the standard of over-all pleasantness, and a con-attitude toward it insofar as it fails to fulfill this standard.

I have said that the *normal* standard of intrinsic value is felt pleasantness in the quality being experienced. There is nothing logically necessary about this. It would be possible in principle (though perhaps not empirically possible for psychological reasons) for a person to have a pro-attitude toward an experience because it was on the whole unpleasant to him and a con-attitude toward an experience which he found on the whole enjoyable. It is logically possible for people to appeal to some other standard of intrinsic evaluation than pleasantness or satisfaction because standards of evaluation are things which people *adopt* (i.e., *decide* to use). It must always make sense to suppose their not adopting a given standard, even if all people do use that standard. This is one respect in which intrinsic value does not differ from extrinsic value.

In light of the foregoing comment we cannot answer the question "What makes a standard of evaluation a standard of intrinsic value?" by saying that the standard is one of pleasantness or satisfaction in the quality of an experience. This particular standard may not be used in an evaluation of intrinsic value. The answer to the question is that one evaluates a quality of experience according to a standard of intrinsic value when one grades it or ranks it regardless of its relation to the value of anything else (except the value of a member of the given class of comparison). *Any* standard may be used as long as it is used to judge a quality of experience without regard to any value it may derive from its relation to other things. The process of evaluation, when concerned with intrinsic value, takes into consideration only the merits and demerits of each quality by itself, entirely disregarding its causes, effects, accompanying conditions, or the wholes of which it forms a part.

It is legitimate to talk of intrinsic value as being "primary" and extrinsic value as being "secondary" only if we mean that extrinsic value is causally dependent upon intrinsic value, while intrinsic value is not causally dependent upon extrinsic value. It is possible for a thing to have extrinsic value because something else has extrinsic value whose value in turn depends on the extrinsic value of a third thing, and so on. But we always arrive finally at something whose value does not depend on the value of anything else, and this is something (a quality of experience) having intrinsic value. The reason for this does not follow necessarily from the definitions of "intrinsic value" and "extrinsic value," since there is nothing inconsistent in supposing a world where intrinsically valuable experiences are not related in any way to extrinsically valuable objects or situations. We could imagine a world where all values were extrinsic, the value of one thing depending on the value of another whose value in turn depends on the value of something else, ad infinitum, although our actual world is not like this. There are causal relations and part-whole relations which hold between intrinsic and extrinsic values, as well as among intrinsic values themselves and among extrinsic values themselves. I shall examine these relations after considering the three kinds of extrinsic value.

5. THE THREE KINDS OF EXTRINSIC VALUE

Inherent Value. The first kind of standard which determines the extrinsic value of something is a standard of inherent value. Inherent value is the capacity of an object, event, situation, process, or any kind of thing *other than* a quality of our own experience, to produce in us when we respond to it (by perceiving it, imagining it, thinking about it, or otherwise apprehending it) a quality of experience which has intrinsic value. Something has inherent disvalue insofar as it can arouse in us a quality of experience which is intrinsically disvaluable. Assuming the normal standard of intrinsic evaluation, which is satisfaction or pleasantness, we grade an evaluatum as inherently good when we find satisfaction in our response to it and as inherently bad when we find our response on the whole unpleasant. The *inherent* goodness of a good thing is its capacity to produce *intrinsic* goodness in the experience of anyone who responds to it. We rank an object according to a standard of inherent value when we compare this capacity of the object with the same capacities of all other objects in the class of comparison. With respect to its inherent value, the good-making characteristics of an object are those which give it the capacity to produce intrinsically valuable experiences in those who respond to it. Similarly, its bad-making characteristics are those which enable it to produce intrinsically disvaluable experiences.

To rank an object as inherently good on the whole is to claim two things—that the object has a greater capacity than most objects in the class of comparison to produce intrinsic value (or a lesser capacity than most to produce intrinsic disvalue); and that the object has the capacity to produce, in a greater number of experiences of a greater number of people, an overbalance of intrinsic goodness rather than of intrinsic badness. Both of these factors are matters which can be determined by straightforward empirical procedures. An example of an evaluation based on a standard of inherent value would be the judgment that a painting is good because, compared to most other paintings in a given class of comparison, it is more frequently pleasant than unpleasant for more people to contemplate aesthetically.

It is quite possible for an object to be inherently good to one person (because it gives him satisfaction) and inherently bad to another (because it gives him dissatisfaction). It is also possible for an object to be inherently good to a person at one time and inherently bad to the same person at another time. An object is inherently good on the whole if it is more often inherently good than inherently bad to those who respond to it.

This definition of "inherent value" implies that there could be no inherent values unless there were intrinsic values, although there could be, in principle, intrinsic values without there being inherent values. It also implies that all inherent values are extrinsic because nothing could have inherent value all by itself. The inherent value of an object depends on the intrinsic value of someone's experience in response to it. If no one ever responded to an object, it could only be said to have *potential* inherent value (or disvalue); if some-one *were* to respond to it he would find satisfaction (or dissatisfaction) in the response. If no quality of anyone's actual or possible experience of the object were intrinsically good or bad, the object could not be said to have even potential inherent value or disvalue.

Instrumental Value. The second kind of standard of extrinsic value is that of instrumental value. In *An Analysis of Knowledge and Valuation* (p. 392), Professor Lewis makes a convenient distinction between instrumental value and utility. A thing has utility to the extent that it is effective in bringing about a given end. If the end is not itself intrinsically or extrinsically good or bad the means to it has no instrumental value or disvalue. Something lacks utility when it is not helpful in bringing about a given end, and a thing has "disutility" when it tends to hinder or prevent the bringing about of a given end. The best means to a given end is that which has the greatest utility, and this is the means which is both sufficient to bring about the end and can bring it about most efficiently. What might be termed the "worst means" is that which is both sufficient to prevent the bringing about of the end and can prevent it most efficiently.

Whether something that has utility also has instrumental value or disvalue depends on whether the given end itself has value or disvalue. (This is what makes all instrumental value extrinsic, i.e., derivative.) There are two ways in which an end can have value and two ways in which it can have disvalue. It has value if it consists in the creation, furtherance, preservation, strengthening, or increase of something which is intrinsically or extrinsically valuable, or if it consists in the destruction, hindrance, avoidance, weakening, or decrease of something which is intrinsically or extrinsically disvaluable. An end can have disvalue in two ways, if it is the creation or strengthening of something disvaluable, or if it is the destruction or weakening of something valuable.

Evaluation according to standards of instrumental value is a

process of grading or ranking something as a means to a valuable or disvaluable end. When the given end has value, the evaluatum is ranked as instrumentally good insofar as it has greater utility than most of the other members of the class of comparison; it is ranked as instrumentally bad insofar as it has greater disutility than most of the other members of the class of comparison. When the given end has disvalue, the contrary judgments are made. In each case a pro-attitude on the part of the evaluator is taken toward a means which is judged to be instrumentally good and a con-attitude is taken toward a means judged instrumentally bad. A neutral attitude is taken toward means which fall at or near the center of the order (having utility—lacking utility—having disutility) in which all members of the class of comparison are arranged. If most of the members of the class of comparison have disutility for a good end, while the evaluatum merely lacks utility, the evaluatum will be preferred and will be the object of a pro-attitude. As another possibility, the evaluatum may lack utility for a good end, while most members of the class of comparison have utility for that end. In that case the evaluatum will be ranked as instrumentally bad (on the whole), and a con-attitude will be taken toward it.

There is an important difference between instrumental value and inherent value. Qualities of experience (which are the only kind of thing that can have intrinsic value) may have instrumental value but not inherent value. When an immediately felt or perceived quality of our own experience has instrumental value, it is a means to something else which has value for us. For example, the aesthetic emotions experienced in a church may be evaluated positively as a means to the deepening or strengthening of one's religious attitudes. Ordinarily the aesthetic emotions will also have intrinsic value, and their instrumental value is simply additional. But we cannot say that such emotions have inherent value, since it makes no sense to say that the felt quality of an emotion has the capacity to arouse a pleasant response in us when we apprehend it. The emotion either is pleasant, and so has intrinsic value, or it is not. It may be capable of arousing further pleasant emotions in us, but then it would take on instrumental, not inherent, value.

Considering the ways in which things can have positive or negative instrumental value in combination with other kinds of value, sixteen different possibilities emerge. When the evaluatum is a quality of experience we have the following combinations:

1. It has intrinsic value, and it has instrumental value as a means to other intrinsic values (or disvalues).

2. It has intrinsic value, and it has instrumental value as a means to extrinsic values (or disvalues).

3. It has intrinsic value, but it has instrumental disvalue as a means to other intrinsic values (or disvalues).

4. It has intrinsic value, but it has instrumental disvalue as a means to extrinsic values (or disvalues).

5-8. Same as 1-4, except that the first clause reads: "It has intrinsic disvalue."

9-16. Same as 1-8, except that the first clause reads: "It has extrinsic value (disvalue) other than the instrumental value (disvalue) mentioned in the second clause."

When the evaluatum is not a quality of experience, only combinations 9-16 hold, since things that are not qualities of experience can only have extrinsic value or disvalue.

Contributive Value. The third kind of standard of extrinsic evaluation is that of contributive value. Something has contributive value if it is a part of a whole which has intrinsic, inherent, or instrumental value. A thing has contributive disvalue if it is a part of a disvaluable whole. Professor Baylis gives the example of a spark plug which has little or no value by itself, but is a valuable thing as part of a gasoline motor (which in turn has instrumental value). Similarly, we could say that a certain series of musical notes have little or no value by themselves but take on contributive value when played as part of a symphony. Indeed, listening to the notes in isolation may have intrinsic disvalue for a person, while listening to the whole symphony has intrinsic value for him. Contributive value is like instrumental value, in that any kind of evaluatum may possess it. The degree to which something has contributive value is in direct proportion to what it contributes to a good whole. If it is a necessary part of the whole (i.e., a part without which the whole would entirely lose its value) or if it is an important part of the whole (i.e., a part without which the whole would lose much of its value), then it is contributively valuable to a high degree. If it is an unnecessary or unimportant part of a good whole, it is graded as less valuable. Corresponding negative judgments occur for a bad whole.

Sixteen possible combinations of contributive value or disvalue with other types of value can be made:

1. Something has intrinsic value, and it has contributive value as part of an intrinsically valuable whole. (An important instance of this, namely the contribution of intrinsically valuable experience to the good life as a whole, is discussed in detail by Professor Lewis in Chapter XVI of *An Analysis of Knowledge and Valuation.* However, in light of what I shall say in Chapter 6 it will be seen that "the good life" can here mean only one way of life among many.)

2. It has intrinsic value, and it has contributive value as part of an extrinsically valuable whole.

3. It has intrinsic value, but it has contributive disvalue as part of an intrinsically disvaluable whole. (That is, the particular experience is enjoyed although it is part of a complex experience which is on the whole unpleasant.)

4. It has intrinsic value, but it has contributive disvalue as part of an extrinsically disvaluable whole.

5-8. Same as 1-4, except that the first clause reads: "It has intrinsic disvalue."

9-16. Same as 1-8, except that the first clause reads: "It has extrinsic value (disvalue) other than the contributive value (disvalue) mentioned in the second clause."

As was true in the case of the sixteen combinations involving instrumental value, when the evaluatum is not a quality of experience only 9-16 hold, since things that are not qualities of experience can only have extrinsic value or disvalue.

This completes my account of the three kinds of standards of extrinsic value. It can now be seen why all extrinsic values in actual life are ultimately dependent upon intrinsic values. The inherent value of something is its capacity for arousing intrinsic values, so by definition this kind of extrinsic value depends on intrinsic value. Instrumental value and contributive value are dependent on intrinsic value whenever the instrumental value of something is due to its being a means to an inherently or intrinsically valuable end, and whenever its contributive value is due to its being a part of an inherently or intrinsically valuable whole. There are two more complex cases, in which instrumental and contributive value are indirectly dependent on intrinsic value. First, a thing may have instrumental value in virtue of the fact that it is a means to an end which has contributive value as part of a whole. This whole in turn either may have inherent or intrinsic value itself or may have instrumental value as a means to an end which has inherent or intrinsic value. Secondly, a thing may have contributive value in virtue of the fact that it is part of a whole which has instrumental value as a means to an end. This end in turn either may have inherent or intrinsic value itself or may have contributive value as part of a whole which has inherent or intrinsic value. Still more complex cases than these can be constructed.

In spite of these connections between instrumental and contributive values on the one hand and intrinsic values on the other, it is possible tò imagine without self-contradiction a world in which instrumental and contributive values did *not* depend on intrinsic value. In such a world, things would be judged as good solely because they were means to ends which were judged to be good solely as parts of wholes which were judged to be good solely as means to ends, and so on indefinitely. It would be a world devoid of inherent and intrinsic value and thus would be very different from our own world. No one would do anything for its own sake, simply because he found personal enjoyment in it. It would be a world of "practical people" who knew how to get things done but had no reason for getting one thing done rather than another. In our world, on the other hand, things are often judged to be good because they make possible, or because they are constituents of, things that we directly enjoy. And sometimes we judge things to be good because *they* are what we directly enjoy. Or we judge our experiences themselves to be good simply because they have the qualities they have. Thus we arrive finally at things which we judge to be desirable in themselves. Their value is intrinsic to them.

C. *Evaluation according to rules*

So far I have been concerned with the process of evaluation when the norms used in it are standards. There is a second basic

type of evaluation in which the norms used are rules. When so used, rules operate as norms for *correct* and *incorrect* behavior or thought. They may be rules governing the use of a language, rules of a game, moral rules, rules of etiquette, laws (legal rules), rules of inference, aesthetic rules (i.e., rules governing what is alleged to be the proper way to create or to appreciate works of art), rules or maxims that guide us in achieving some personal goal (such as health or wealth), rules or regulations which govern a social institution (a club, a political party, a labor union, a university, etc.), or rules which govern practices or procedures (parliamentary procedure, business office procedures, procedures in a court of law, etc.). The typical words which are used to express evaluations according to rules are "right" and "wrong." The typical words used in evaluations according to standards are as we have seen, "good" and "bad." (I shall consider other ways of expressing value judgments in Chapter 2.)

All evaluations according to rules are gradings; none are rankings. This is because there are no degrees of right and wrong as there are degrees of goodness and badness. I have already pointed out that something may fulfill or fail to fulfill a *standard* of evaluation to a certain degree. But an act either complies with or breaks a rule. If we refer two acts to a rule which applies to both, either we find both acts equally right (they are in accordance with the rule), or both equally wrong (they violate the rule), or one right and the other wrong. Of course a right act is better than a wrong act and a wrong act is worse than a right one, but this hardly deserves to be called "ranking" the two acts. What determines degrees of goodness and badness is the extent to which the evaluatum fulfills or fails to fulfill a standard. What determines whether the evaluatum is right or wrong is whether or not it is in accordance with a rule that applies to it. It is to be noted that an act may be neither right nor wrong, relative to a given rule. This would occur when the rule does not "cover" the act, that is, the act is not included in its range of application. The rule "Automobiles are forbidden to park on the bridge" does not apply to a person walking or riding a bicycle across the bridge.

Rules are of three sorts—injunctions or requirements, permissions or options, and prohibitions or forbiddings. That is to say, when a rule includes an act in its range of application the act may be enjoined by the rule, permitted by the rule, or forbidden by the rule. An example of the first is, "One must have a license to drive a car"; of the second, "Smoking is permitted in this room"; of the third, "No smoking in this room." Sometimes a "right act" is defined as one that is *obligatory*, that is, as one which complies with a rule that enjoins or requires it to be done. (Moral obligation is a special case of this.) At other times it is defined as any act that is not wrong, and a wrong act is defined as an act prohibited or forbidden by a rule. In this second sense, any act which is *either obligatory or permissible* is a right act. I shall not use the word "right" in either of these ways. By saying that an act is right I shall mean simply that it complies with a rule that covers it. Thus an act is right in my sense under any of three conditions: (1) It is in accordance with a rule that requires it to be done. (2) It is in accordance with a rule that permits it to be done. (3) It is in accordance with a rule that forbids it to be done. The third condition may be stated less awkwardly in the following way: When a rule forbids an act to be done, doing the act is wrong and refraining from doing the act is right. An act is wrong in my sense whenever it does not comply with a rule that covers it. Thus refraining from doing an obligatory act is wrong and doing a prohibited act is wrong. No wrong act corresponds to the second type of rule, since the rule permits the doing of an act without forbidding or requiring the doing of any other act. If one does not do the permitted act, one is not breaking the rule and hence no wrong is committed. In these cases I shall say that doing the act is optional, meaning that it is not wrong to do it and that it is not wrong to refrain from doing it. Both alternatives comply with the rule and therefore both are right. In contrast, I shall restrict the phrase "neither right nor wrong" to acts *not covered* by a rule.

Rules of the first and third types (i.e., injunctions and prohibitions) are related in the following way. To require or enjoin the performance of an act is to forbid or prohibit its nonperformance; it is wrong to *refrain* from doing a required act and right to *refrain* from doing a forbidden act. And it is right to *do* a required act and wrong to *do* a forbidden act. It is also possible to define rules of the first and third types as negations of permissions; a rule which

requires that an act be done does not permit its not being done. In other words, refraining from doing a required act is not permitted and consequently the rule requiring the performance of the act may be thought of as not permitting its nonperformance. Similarly a rule of the third type, which forbids that an act be done, does not permit its being done. It is a rule not permitting the performance of an act, in contrast to a rule of the first type, which is a rule not permitting the nonperformance of an act.

Professor Bernard Mayo has argued that because there are no degrees of rule-following and rule-breaking and accordingly no degrees of right and wrong, judgments of right and wrong are not evaluations. Such judgments, he says,

> . . . are not evaluative, but imperative. To say that an action is wrong is not to apply a scale of preference, since it is not to say that there are other actions which are less wrong. An action is either right or wrong or neutral: it cannot be more or less wrong, as it can be more or less evil. An action is right if it conforms to an accepted principle or rule, wrong if it violates one. Since there are no degrees of rule-breaking—a rule is either observed or it is infringed—there can be no degrees of right and wrong. . . . The logic of the application of principles is thus different from the logic of the application of standards. Standards are fallen short of by greater or lesser degrees, but principles are either conformed to or not. (B. Mayo, *Ethics and the Moral Life*, London: Macmillan & Co., Ltd., 1958, p. 29.)

I would not deny anything that is said here about the use of the terms "right" and "wrong." But I would classify judgments of right and wrong as evaluations, not imperatives (i.e., prescriptives).

Judgments according to rules are similar in certain basic respects to judgments of goodness and badness, the very factors which set the logical pattern of the process of evaluation being common to both. Whether a judgment be concerned with right or wrong, or with goodness or badness, it is the outcome of a process of evaluation in which the following five factors are present: (1) a class of comparison; (2) a norm or set of norms according to which the evaluatum is judged; (3) a pro-attitude, a con-attitude, or a neutral attitude; (4) a set of good-making and bad-making characteristics (now to be referred to as right-making and wrong-making characteristics); and (5) a point of view from which the evaluation is made. Let us examine the place of each of these elements in an evaluation according to rules.

Class of Comparison. In evaluation according to rules, the sorts of things that may serve as members of a class of comparison are the sorts of things that can be judged right or wrong. There are three of these: human decisions, choices, and acts. We speak of making a right or wrong decision or choice, as well as of doing a right or wrong act.

I have already shown how rules can enjoin, permit, or prohibit acts, and hence how acts can be judged according to rules. But do rules enjoin, permit, or prohibit decisions? They do so indirectly. Deciding may be deciding to do something or deciding that something is the case. In both instances we resolve our doubts about something; we make up our mind either to do one thing rather than another, or to claim that something is or is not the case. And we decide either rightly or wrongly, depending on whether the act we decide to do is right or wrong, or on whether the claim we decide to make is justifiable or unjustifiable. Rules of conduct (indirectly) guide the first kind of decision, rules of thought (indirectly) guide the second. A decision is right, then, if it is a decision to do a right act (i.e., a decision to act in accordance with a rule of conduct) or if it is a decision to make a justifiable claim (i.e., a decision to assert or deny something in accordance with rules of valid reasoning). To say "You made the right decision" is to say either "You decided to do the right thing" or "You decided correctly to assert (deny) that such-and-such is so." It is in this way that decisions can be judged according to rules. They can also, of course, be judged according to standards. Thus there may be good or bad decisions as well as right or wrong decisions.

Just as rules may function as norms for decision-making, so they may also function as norms for choice-making. "Choice" has a double meaning in ordinary language. Sometimes it means selecting-one-thing-rather-than-another. In this sense, to choose is to do something, since the selection is an act. Examples would be choosing a lot on which to build a house or choosing a dessert from a list on a menu. In these cases the act of choosing normally consists in a linguistic act, such as saying "I'll take the lot on the corner"

or "I'll have apple pie." If a person holds out a pack of cards to us and says "Choose a card," our choice may consist in the nonlinguistic act of taking a card. Since they are acts, choices in this sense may be evaluated as right or wrong choices according to rules that enjoin, permit, or prohibit them.

In a second sense choices are not acts but are either similar to or identical with decisions. Choices are similar to decisions in that a person may choose to do something and then not actually do it (just as a person may decide to do something and then not carry out his decision). A choice in this sense is a preference for doing one thing rather than another. "My choice was to go to the other restaurant but my friends all wanted to eat here." Choosing is not identical with deciding, however, unless one's choice is the outcome of deliberation which resolves a doubt in one's mind. One may choose to do something merely on the basis of one's emotions or moods. Deciding on the other hand contextually implies a process of deliberation; choosing has no such implication. We can say that all decisions are choices (in this second sense of "choice"), but we cannot say that all choices are decisions. Choices can be spontaneous or thoughtless, decisions cannot. Unlike decisions, choices can be based solely on our immediate likes and dislikes. But whether choices are deliberative or spontaneous, they may be evaluated according to rules which (indirectly) apply to them. We have seen that choices in the sense of acts are subject to rules *directly*. In these cases to say "You made the right choice" means "You did the right thing in choosing the way you did." In the second sense, to say "You made the right choice" means "You chose to do the right thing." As in the case of decisions, rules are indirectly applicable here. As a right decision is a decision to do what is right, so a right choice is a choice of (preference for) what is right.

For the sake of simplicity, in all subsequent discussion of evaluation according to rules, I shall refer to acts rather than to decisions and choices when speaking of the evaluatum or the members of a class of comparison. That is, I shall only examine cases in which rules *directly* apply to an evaluatum or a member of a class of comparison.

Appeal to Norms. When we evaluate something according to rules, the rules function as norms in the same way as do standards of evaluation. The process of evaluation according to rules yields a value judgment in which we claim that an act is right or wrong. In order to arrive at such a judgment, we must determine two things: (1) whether a given act falls under (is covered by) a rule and (2) if it does fall under the rule, whether doing the act is in accordance with or violates the rule. If the rule enjoins or requires the act, doing the act is in accordance with the rule and not doing it is a violation of the rule. If the rule permits the act, either doing it or refraining from it is in accordance with the rule. If the rule forbids the act, then doing it violates the rule and refraining from it is in accordance with the rule. A rule includes an act in its range of application (and the act is "covered by" or "falls under" the rule) when the act is of a sort specified in the rule and the performance or nonperformance of the act occurs in circumstances specified in the rule. Thus the rule "No smoking in subways" covers all acts of smoking and all nonperformances of the act of smoking (i.e., all cases of refraining from smoking) by persons while they are in subways. It does not cover acts of smoking done (or not done) by persons outside of subways. To take another instance, the rule "Promises must not be broken" applies only to people who have made promises and only to those acts which could correctly be construed as the keeping or the breaking of promises made.

Pro-Attitudes and Con-Attitudes. The class of acts covered by a rule is the class of comparison for all evaluations made according to the rule. Any act that falls under a rule may be *graded* as right or as wrong, depending on whether the act is in accordance with or violates the rule. When we judge an act to be right according to a rule, we have adopted the rule as the basis (norm) for judging it. In doing so, we commit ourselves to having a pro-attitude toward any act complying with the rule and a con-attitude toward any act violating it. We have a pro-attitude toward a person's doing a required act (or refraining from a forbidden act), and we have a con-attitude toward a person's not doing a required act (or doing a forbidden act). When our rule is a permission we take a neutral attitude toward any act in the specified circumstances, whether it be performance or nonperformance. With regard to acts not covered by a rule we have adopted, we have *no* attitude, not even a

neutral one. So far as our commitment goes, such acts are indifferent; they are neither right nor wrong. They become right or wrong, however, the moment we adopt a rule which includes them in its range of application. Thus when a given rule demands, permits, or forbids an act to be done in certain circumstances, the class of comparison is the class of all acts that are correctly describable either as the performance of what is demanded, permitted, or forbidden in the specified circumstances or as the nonperformance of what is demanded, permitted, or forbidden in those circumstances. The rule serves as a norm for judging or grading a given number of the class of comparison (namely the evaluatum). If the evaluatum fulfills the norm, the judgment is positive (or neutral, in the case of permissions); the evaluator accordingly takes a pro-attitude (or neutral attitude) toward the evaluatum. A con-attitude is taken toward the evaluatum when it fails to fulfill the norm, that is, when it violates an injunction or prohibition.

It has been pointed out that there are no degrees of right and wrong in the way that there are degrees of goodness and badness. In the case of evaluations according to standards, an evaluator's pro-attitude or con-attitude varies in strength according to the degree of goodness or badness of the evaluatum. What determines the strength of pro- and con-attitudes in evaluations according to rules? The answer lies in the degree of *stringency* or *importance* of the rule on which the evaluation is based. The more stringent the rule, the stronger is the evaluator's pro-attitude toward an act done in accordance with it, and the stronger is his con-attitude toward a violation of it. The degree of stringency or importance of a rule is determined by an evaluation of the rule according to standards of extrinsic value. A rule itself may function as an evaluatum and be judged according to standards; it may have inherent, instrumental, or contributive value (or disvalue), depending on what standards are used in grading or ranking it. For example, if we judge a rule according to the utilitarian standard, we determine its instrumental value (or disvalue) as a means to the general happiness of mankind. Under such an evaluation the rule "Promises are not to be broken" is seen to be much more stringent or important than the rule "Brush your teeth twice a day." In morals the classic difference between duties of obligation ("perfect" duties) and duties of supererogation ("imperfect" duties) may be understood in this way, the first consisting in acts either enjoined or forbidden by highly stringent moral rules, the second consisting in acts enjoined, permitted, or forbidden by less stringent moral rules. (However, some philosophers would make this a difference not of degree but of kind. Perfect duties would be *right* acts enjoined by moral rules; imperfect duties would be *good* acts which fulfill moral standards.)

Right-making and Wrong-making Characteristics. It is easy to see how right-making and wrong-making characteristics function in evaluations according to rules. A right-making characteristic of an act that falls under a rule is any feature of the act in virtue of which it complies with the rule. The wrong-making characteristics of an act are those features which make it a violation of the rule. Can an act have both right-making and wrong-making characteristics with respect to the same rule? The answer depends on two factors: how complex the act is and how well formed the rule is. A businessman's complex act of carrying on an advertizing campaign for his product can have both right-making and wrong-making characteristics with regard to the vague rule "Be honest." But suppose we adopt a clearly formulated rule, such as: "Whenever a statement made in a public medium of communication will be taken by others as a statement of fact, it must not be a false statement and it must be free from distortion and exaggeration." If we apply this rule to a single statement made in one of the advertisements in the campaign, the act of making the statement will have only right-making or wrong-making characteristics.

Point of View. The final element common to evaluations according to both standards and rules is the presence of a point of view from which the evaluation is carried out. Just as we may judge a painting from an aesthetic point of view, an economic point of view, and a moral point of view, so we may judge an act from many different points of view, only using as our norm of evaluation rules and not standards. We may judge the same act to be right from one point of view (applying its appropriate rules) and wrong from another point of view (applying *its* appropriate rules). For example, the rules of good business practice might require a businessman in times of recession to discharge a number of his employees; such an act would be right from the business point of view. But it might

be wrong from the moral point of view, if other economizing measures could be used instead. Similarly, an act which was right from the aesthetic point of view, such as constructing a fine work of architecture as a city's post office, might be wrong from a political or economic point of view. In all cases of this sort an act has both right-making and wrong-making characteristics, but these are determined by different rules. The act is both right and wrong, but right economically and wrong morally, or right aesthetically and wrong politically. Whether the act is right or wrong *on the whole* depends on which rules take precedence over others. (I shall consider the relative precedence of rules along with the relative precedence of standards in Chapter 3.)

I conclude, then, that judgments of right and wrong are sufficiently similar in important respects to judgments of goodness and badness to warrant their being included in the concept of evaluation. The parallel between evaluations according to rules and evaluations according to standards will continue to be drawn in subsequent chapters. That standards and rules have essentially the same function in evaluations is a principal thesis of this book, and I shall also try to show that standards and rules are both subject to the same methods of reasoning when we try to justify our adopting them in practical life.

D. Deliberation

My purpose in the remainder of this chapter is to make clear the relation between the concept of evaluation and the concept of deliberation. There are two sorts of deliberation: deliberation about acts and deliberation about standards and rules. This distinction corresponds to the one which Mr. R. M. Hare has drawn between making decisions, on the basis of given principles, about what we shall do and making "decisions of principle," that is, deciding what principles to adopt as guides to our future decisions. In his discussion of this distinction (in Chapter 4 of *The Language of Morals*. Oxford: The Clarendon Press, 1952) Mr. Hare uses the example of learning to drive a car. Some of our learning process consists in being taught principles of good driving. Having learned these principles, we know what to do to start the car, stop the car, turn around, speed up, slow down, and so forth. The principles guide our decisions about what to do in a certain situation. Unless we had made a decision to adopt these principles as guides to our conduct, they would be no help in learning how to drive. Indeed, we could not learn how to drive at all (or how to do anything else, for that matter) unless we *decided* to use principles as guides to our decisions. But learning how to drive does not consist solely in adopting principles and applying them in practical situations. We are never taught enough principles to guide us in every situation we will meet when we are driving. As we encounter more and more various circumstances, we have to make up our own principles. We set precedents for ourselves by imposing new principles upon our own conduct. We also learn in this manner when to make exceptions to the principles that have been taught us.

> Suppose that we have a principle to act in a certain way in certain circumstances. Suppose then that we find ourselves in circumstances which fall under the principle, but which have certain other peculiar features, not met before, which make us ask 'Is the principle really intended to cover cases like this, or is it incompletely specified—is there here a case belonging to a class which should be treated as exceptional?' Our answer to this question will be a decision, but a decision of principle, as is shown by the use of the value-word 'should.' If we decide that this should be an exception, we thereby modify the principle by laying down an exception to it. (R. M. Hare, *op. cit.*, p. 65.)

By functioning as guides to our conduct, principles serve as norms by which we learn what is the right thing to do and what is the wrong thing to do. To break a principle is wrong, unless we can justify making an exception to it or we can justify abandoning it and adopting a new one. In this way we come to judge our conduct and our decisions according to our principles. Such judgments are value judgments, and the principles according to which they are made are what I have called standards and rules. Whether a given decision or act will be a good one or a bad one, a right one or a wrong one, depends on the standard or rule we have adopted as the ground for our evaluation.

Corresponding to this distinction between decisions based on already adopted principles and decisions to adopt principles is a

distinction between two kinds of deliberation. In the first kind of deliberation we are trying to make a decision of the first sort. That is, we are trying to decide what to do, what course of action to take from among the alternatives confronting us. We are not questioning the standards and rules which we think are applicable to the situation. In the second kind of deliberation we are uncertain about what standards and rules to apply. We are trying to decide, in Mr. Hare's terms, what principles to adopt, or perhaps whether to make an exception to a principle. I shall consider each kind of deliberation in turn.

When, in a situation of choice, we are deliberating about alternative courses of action, we are trying to decide which course of action is the best. In order to do this we must carry out a series of evaluations in which we *grade* each of the various alternatives and then *rank* all of them in a preferential order from best to worst. Let us consider these two steps of the process of deliberation in detail. (It should be remembered that, like the foregoing account of the process of evaluation, this is not intended to be a psychological description of the way people actually think, but an explication of the logic of their thinking when it is ideally rational.)

1. We carry out an evaluation of each alternative act according to *rules* we have accepted as appropriate or valid in the situation at hand. Our purpose here is to discover which acts are right and which are wrong, as judged according to all the rules we think the acts fall under. If there are rules that cover all the alternatives and if only one alternative can be graded as a right act according to these rules (i.e., all the other acts violate the rules and are therefore wrong), then our deliberation is completed. We know that the one alternative which complies with the rules is not only *a* right act; we also know that it is *the* right act. It is, then, the best thing to do in the circumstances. If on the other hand either our rules do not cover all the alternatives (so that some of them are neither right nor wrong) or more than one act is found to be right, then we must proceed to the second step of deliberation.

2. We carry out an evaluation of each alternative according to *standards* we have accepted, to see which alternative is the most valuable or desirable. The class of comparison of such evaluations is precisely the class of alternative courses of action open to the

evaluator's choice. The aim of this step is to arrange the entire set of alternatives (i.e., the members of the class of comparison) in an order of increasing value or desirability. The act which has the highest place in this order will be the best thing to do, that is, the course of action we ought to choose. Usually this alternative will be an act that was found to be right in our first step of deliberation. But it need not be. If all the acts, when judged according to rules, are wrong, there may still be one which is less *bad* (i.e., whose effects are less undesirable) than the others. Similarly if all the acts are right acts according to our rules, some will be worse than others and one will be better than all the others, when judged according to our standards.

The difference between *a* right thing to do and *the* right thing to do and the difference between *a* wrong thing to do and *the* wrong thing to do may be made clear by reference to the situations in which, respectively, all the alternatives are right acts and all are wrong acts. When all the alternatives are right acts, as judged by our rules, then every act is *a* right thing to do. But only one of the acts is *the* right thing to do, that is, the thing we ought to do. This is the act that is better than any other alternative, when judged according to our standards. This act is the one which not only has value (goodness) in addition to its rightness, but has greater value than any of the other acts. This situation also allows us to distinguish between a required or enjoined act and an obligatory act. (This distinction was disregarded in Section C of this chapter, where I spoke of an obligatory act as a required or enjoined act.) A required or enjoined act is any act which is required or enjoined by a rule. Now it is possible in principle for all the acts in a situation of choice to be enjoined acts. But only one act can be obligatory. That is to say, a person can be obligated to do only one act when he must choose among several. He cannot be obligated to do what is impossible, and presumably it is impossible for him in the given situation to do more than one of the alternatives. An obligatory act, then, is an enjoined act which is *the* right thing to do in a situation of choice. It is the act which is better than any of the alternatives.

With regard to the second case, where all the alternatives are wrong acts, we are tempted to say that the best of these acts is

the one which is least wrong. But this is to confuse "least wrong" with "least evil (or bad)." In such a situation our choice is a choice of a lesser evil, not of a lesser wrong. There is no lesser wrong. All the acts are equally wrong, although some are worse (less good or more evil) than others. The best act is that one which, when compared with the others, has the greatest value or the least disvalue according to our standards. Now even though such an act is *a* wrong thing to do, it is not *the* wrong thing to do. In fact in the given situation it is *the* right thing to do (although it is not *a* right thing to do). Every one of the acts in the given situation is *a* wrong thing to do and *all but one* is *the* wrong thing to do. *The* wrong thing to do is any alternative except the best, which is *the* right thing to do. But what is *the* right thing to do in this situation is *a* wrong thing to do. (Sometimes we call a situation of this kind "tragic.")

By means of the second step of deliberation, then, we decide what we ought to do whenever an evaluation of alternatives according to rules yields more than one right act or no right acts. This situation will also occur whenever we find by means of the first step that all the alternatives are permissible and none is enjoined. So far as our rules are concerned, we may do anything we like. Nothing we do will be wrong. But there is nevertheless one act which is *the* right thing to do, and that is simply the *best* thing to do as judged by our standards. We must then enter upon the second step of deliberation in order to arrive at a rational decision about what we ought to do.

Let us now turn to the second sort of deliberation. Here we are not trying to decide what to do in light of our rules and standards; we are trying to decide what rules and standards to adopt. We must evaluate rules and standards according to "higher" standards, that is, according to standards that take precedence over the rules and standards to be evaluated. The class of comparison is the entire group of rules and standards from which we are trying to choose the most valid or appropriate for our evaluations. The overall purpose of our deliberation is to find the best rules and standards to adopt, not only as guides to our decisions and conduct, but also as the norms for our evaluations and prescriptions.

A standard used for evaluating standards and rules might be called a "second-order" standard. When we evaluate rules according to second-order standards, we judge the *effects* of their being followed. We consider what it would be like if one rule rather than another were generally adopted as a norm for right conduct. We then try to determine to what extent the probable consequences of its adoption would fulfill or fail to fulfill the second-order standards. In ethics this system is sometimes called "restricted utilitarianism." Such a system, however, need not operate only for moral rules. A businessman, the head of a corporation, might try in similar fashion to find the best rules for running his company with maximum economy and efficiency. Again, people might decide to change certain rules of etiquette because they have become inconvenient or impracticable under changing conditions of social life. And, of course, legislators may use second-order standards when deliberating about which new laws to enact and which old laws to change or abolish. (Laws are one kind of rules.) The particular second-order standards people will actually use are the standards they consider *appropriate* or *valid* for evaluating different sorts of rules governing different sorts of conduct in different sorts of circumstances. (The grounds for determining the appropriateness or validity of second-order standards will be discussed in Chapter 3.)

When we evaluate *standards* rather than rules, a similar procedure of evaluation is used. We determine what practical consequences are likely to be entailed if this standard or that were generally adopted as a guide to the decisions and conduct of people. We then consider these consequences in the light of our second-order standards, ranking the group of first-order standards in an order of preferability according to the degree to which each one fulfills or fails to fulfill the second-order standards. For example, suppose that in the past our standards for choosing a family car have been its appearance and power. We might come to adopt new standards (say, efficiency of performance and compactness of size) when we find that the cars we bought in the past have not been satisfactory for family use. Or to take another example, a department store buyer might have been using certain standards for buying dresses which did not sell well. He might then change his standards (or the head of the dress department might insist that he change them!) in the hope that dresses selected on the basis of the new standards would appeal to the store's customers.

In both of these examples the unsatisfactory effects of using the old standards constitute their bad-making characteristics. These bad-making characteristics are determined by the second-order standards that are deemed appropriate or valid for judging the old standards. In the second example the buyer's old standards may have been the expression of his personal taste concerning good fabrics and dress designs. On the basis of the second-order standard (increasing the store's sales), the buyer changes his standards for selecting dresses from what he personally likes to what the customers like. The appropriateness or validity of his second-order standard depends on the conditions under which the evaluation of standards takes place, including the purposes for which the evaluation is carried out. Further consideration of this point will be taken up in Chapter 3. We shall find that the utilitarian methods I have mentioned in the preceding two paragraphs are not the only methods for evaluating rules and standards.

To sum up, all deliberating is evaluating. We deliberate on two levels: we try to decide which course of action is best among a set of alternatives, and we try to determine which standards and rules to adopt as guides to our conduct. In the first kind of deliberation our evaluating is done according to first-order standards or rules; in the second kind, according to second-order standards. The two kinds of deliberation differ, then, with respect to their classes of comparison and with respect to the norms which they use. Apart from this, they are similar to each other and to all other processes of evaluation.