

(morally) justifiable to utter a false value judgment. In support of this he cites the following cases:

It may be true that one's neighbour is acting immorally, but wrong to publish the fact to the whole neighbourhood. It may be necessary to calm a hysterical delinquent by assuring him for a time, that he has done nothing wrong. One may have performed a noble action but it be wrong to remind oneself of the fact. In other words, the question of whether one ought to express an opinion is different from the question of whether that which one might express is true or false. . . . Distinctions of the kind suggested here are often made in ordinary discourse. For example, it is often maintained that one ought not to express an opinion about the conduct of other people. This is not only because it is so difficult to know enough about other people to form a correct opinion. What is maintained is that we have no right to express such opinions, even when they are correct. (E. Gilman, "The Distinctive Purpose of Moral Judgments," *Mind*, LXI, 248, 1952, pp. 311-312.)

If such cases do show that in some circumstances it is wrong to utter a justified value judgment and right to utter an unjustified one, then the question of whether a given judgment is justified is independent of the question of whether the act of pronouncing it (either to others or to oneself) is justified.

The argument becomes even stronger when we consider pragmatic rather than moral justification of linguistic acts. There are many ways in which expressing unjustified value judgments can be an effective means to an individual's or group's ends (and so be pragmatically justified). An unscrupulous art dealer might succeed in selling a painting which he knows to be a fake by praising it highly in the presence of a customer. A candidate for political office might further his ends by unjustifiably defaming his opponent. The government of a country might find it useful to make exaggerated and unwarranted claims about injustices in another country. On the other hand, pronouncing a true value judgment might well work against a person's ends. A selfish man would certainly frustrate his own desires by expressing publicly an honest evaluation of his character.

In order to justify a value judgment, then, it is not sufficient to justify the act of uttering it. We tend to overlook this when we think that we can justify the judgment that it is wrong to steal, for

The justification of value judgments: verification and validation

A. What does it mean to justify a judgment?

Up to this point I have tried to make clear what it is to carry out an evaluation of something and what it is to make a judgment of the value of something. I wish now to consider the way we reason for or against value judgments when the reasons which we give are good (legitimate, sound, warranted, valid, intellectually acceptable) reasons.

It is important at the outset to distinguish between justifying a value judgment and justifying our uttering a value judgment. We justify a particular value judgment when we give good reasons for a person's grading or ranking an object in a certain way. But such reasons are not necessarily good reasons for uttering that particular judgment in a given set of circumstances. It might at first be thought that if a value judgment has been justified, then it is always justifiable to utter it. That this is not the case has been argued by Professor Eric Gilman. He states that it is sometimes (morally) unjustifiable to utter a true value judgment and sometimes

example, by showing that it furthers the ends of morality to say to a child in a certain tone of voice "It is wrong to steal." But from the cases already cited it is clear that the moral goodness or rightness of an act of pronouncing a judgment has nothing to do with whether there are good reasons for accepting the judgment.

When I speak of justifying value judgments, how is the word "justify" being used? It is correct to say that we do not justify propositions but verify them, and that we do not verify decisions, acts, or dispositions, but justify them. We verify (or confirm) a proposition by showing that it is true, that is, by giving the evidence for it, or by offering reasons in support of it, or by proving it, or by specifying the grounds on which it rests and showing that they are good grounds. We justify a decision, an act, or a disposition by giving reasons for *making* the decision, for *doing* the act, or for *having* the disposition. Or else, if there are good reasons both for and against it, we justify making the decision, doing the act, or having the disposition by showing that the reasons for it outweigh the reasons against it. Thus we speak of justifying (not verifying) one's decision to join the army, to buy a new car, or to follow someone's advice. We say that a person is justified in doing a certain act. ("He is justified in firing that employee.") And we speak of a person's attitudes as being justified. ("He is justified in disapproving of that new law.")

Confusion occurs because we talk (idiomatically) about *assertions* and *beliefs* as being verified as well as being justified. We say "His assertion that there is life on Mars can be verified" or "His belief that there is life on Mars can be verified." But we also say "He is justified in making the assertion that there is life on Mars" or "He is justified in believing that there is life on Mars." Since making an assertion is an act (of uttering a declarative sentence under certain conditions) and since believing is a mental disposition, it is legitimate to use the word "justify" in this way. When we speak of verifying an assertion or belief, however, we are not referring to the *act* of asserting or the *disposition* of believing. We are instead referring to *what* is asserted or believed. We are saying that what is asserted or believed can be shown to be true. What we assert is a proposition (namely that something is or is not so), and the same proposition can be asserted in many different ways (i.e., by means

of many different sentences). The proposition is what is true or false, not the act of asserting it. Similarly it is a proposition that is believed, and to believe it may or may not be justified.

There is no logical connection between the verification of a proposition and the justification of an assertion or belief. "That p is true" does not entail, and is not entailed by, "It is justified to assert that p " or "It is justified to believe that p ." Let us first consider cases of asserting that p .

A person is justified in asserting that p only when there are good reasons for him to utter the sentence " p " (or some equivalent sentence). These reasons may be good reasons from the moral point of view, from the point of view of prudence, from the point of view of etiquette, or from the logical or intellectual point of view. That is to say, the person's asserting that p may be justified in any one of these ways, even when p is false. Similarly, a person may not be justified in any of these ways in asserting that p , even when p is true.

Examples may be found, taking each of these points of view. Someone who has promised to keep a secret is not morally justified in making the secret known, even if what he makes known is true. On the other hand, a doctor may be morally justified in making false assertions to a patient whose chances for recovery would be greatly diminished if he were told the truth about his condition. From the point of view of prudence, it may be to a man's self-interest to assert what he knows to be false (as in lying to someone about his accomplishments) and it may go against his self-interest to tell the truth. From the point of view of etiquette, it may be impolite to make a true remark in front of someone, and sometimes etiquette demands that we assert what is false (e.g., saying that we enjoyed ourselves at a party when in fact we did not). There is nothing puzzling about these cases. But how is it possible from the *logical* point of view to be justified in saying what is false and not justified in saying what is true? The answer becomes clear when we realize that a person may assert what is false because he has been misled by the evidence, and that a person may assert what is true as a result of a lucky guess. A jury, for example, is logically (as well as morally) justified in giving a verdict of "not guilty" when there is insufficient evidence for the guilt of the accused or when the circumstantial

evidence points to his innocence, even if he did in fact commit the crime. And a fortune teller who used tea leaves to make predictions is not logically justified in telling a woman that she will marry a dark, handsome man, even if this turns out to be the case.

Turning from asserting that *p* to believing that *p*, I shall only consider cases of being justified or unjustified from the logical point of view. Here again we have two possibilities: a person may not be logically justified in believing what is true, and may be logically justified in believing what is false. A person is not logically justified in believing a true proposition when he is not in a position to *know* that it is true, that is, to have good reasons for believing it to be true. A man who believes that a certain horse will win a race is not justified in believing it if his belief is based on a hunch, or on the fact that the first letter of the horse's name is the same as the first letter of his own name. His belief might well be true; the horse might in fact win the race. But he is justified in believing this only if he has *good* reasons to think it will win, such as reliable knowledge that the race has been fixed. Circumstances in which a person is justified in believing a false proposition are more unusual, but they do occur. A child who has always been told by his parents in a serious tone of voice that there is a Santa Claus is justified in believing that there is a Santa Claus (assuming his parents generally tell him the truth when they speak in a serious tone of voice).

Let us now turn to the use of the word "justify" in reference to value judgments. I have said that there are three kinds of things which we can justify: decisions, acts, and dispositions. (I have just been considering one sort of act, that of making an assertion, and one sort of disposition, that of belief.) A value judgment may be thought of either as a decision or as a disposition. As a decision it is a mental act. (It differs, of course, from behavioral or public acts.) Since it is a mental act of grading or ranking something as a result of a process of evaluation, it is not to be identified with the act of pronouncing or expressing a judgment. We have just seen that it may be justified when the act of uttering it is not justified, and vice versa. But what does it mean to say that a value judgment itself is justified? To say that the judgment ascribing a certain value *V* to something *X* is justified, is to say that either the *decision* or the *disposition* to make such a judgment is justified. It is to claim that

the evaluator or judge has good reasons for *deciding* that *X* has *V* or for *believing* that *X* has *V*; or that, if he has good reasons for and also good reasons against such a decision or disposition, the reasons for outweigh the reasons against.

Can value judgments be said to be true or false as well as justified or unjustified? It might be thought that there is the same relation between a value judgment and its content as there is between a belief and a proposition. Indeed, I have just used the phrase "believing that *X* has *V*" to refer to a value judgment as a disposition. Why not draw a parallel here and speak of a judgment (a decision or belief-disposition) as being justified or unjustified, and the content of a judgment (*what* is decided or believed) as being true or false, verified or falsified? We are tempted to speak in this way for two reasons. One is the fact that value judgments are expressible in declarative sentences. The other is the fact that "true" and "false" are often used in ordinary language simply as words of assent and dissent. I shall consider each of these reasons in turn.

When we apply the words "true" and "false" to nonvalue assertions, meaning thereby that what is being asserted is or is not intellectually acceptable, we contextually imply a method of verifying or falsifying the assertions. (Other meanings of the words "true" and "false" will be considered shortly.) When we ask "Is what he said true?" we expect that anyone who answers definitely in the affirmative or negative can show that what he said is or is not true (or at least can point to someone who can show this). Otherwise we do not think one is justified in giving a definite answer to our question. Now it is perfectly idiomatic to use the same phrasing about the utterance of a value judgment in a declarative sentence. When a person says "Object *X* has value *V*" we can ask whether what he says is true. What we mean by our question is, "Does the object *X* really have value *V*, as he says?" And this appears contextually to imply that anyone who answers our question definitely one way or the other can verify or falsify the statement, that is, can show that *X* does have *V* or does not have *V* (or can point to someone who can show this). In both cases we appear to be concerned with a correspondence between what a person says (assets, claims) to be the case and what is the case. If there is a correspondence then what he says is true; if not, then it is false. The conclusion drawn

from this parallel is that value judgments are factual or assertive. As a result, the only difference between value and nonvalue assertions will lie in the *procedures* of verification and falsification. We do not verify or falsify the two sorts of assertion in the same way. But this, it is said, is not a difference in their logical grammar. In both cases words are being used to make an assertion that something is or is not so; in one case we are asserting that an object has value or disvalue, in the other we are asserting something else about an object. In both cases, the argument runs, we can distinguish what is being asserted (the proposition) from the act of asserting it or from the disposition to assert it. We distinguish between the content of a value judgment and the act of judging or the disposition to judge. This, it is claimed, is correlative with the distinction in non-evaluative assertions between what is believed and believing it. In short, the conclusion is drawn that the content of a value judgment can be verified or falsified just as a nonvaluative proposition can be verified or falsified.

What is wrong with this argument? There is no fallacy in it; it is rather a matter of a mistaken emphasis that can be seriously misleading. These similarities between evaluative and non-evaluative assertions tend to make us overlook the differences between them. Once we begin to think of value judgments as a special kind of factual assertion, certain questions seem inevitably to arise. What property is attributed to an object or act when it is judged to be good or right? How can this property be known? Granted that it is not perceivable, can it be inferred from sense perceptions? What is it in the real world to which a value judgment must correspond in order to be true? Is there a realm of values as well as a realm of facts? When philosophers discuss value judgments in this way, they become entangled in the fruitless disputes carried on by intuitionists, naturalists, and nonnaturalists. Their mistake is to have overlooked certain crucial differences between value judgments and nonvaluative assertions. These differences are precisely those which make it misleading to speak of a value judgment as "true" or "false" and consequently as "verified" or "falsified." We shall see that *one* step in the justification of a value judgment is empirical verification. But because there are other essential steps, it is in the interest of clarity

to speak of value judgments as justified or unjustified rather than as true or false.

The second reason why philosophers have thought it proper to apply the terms "true" and "false" to value judgments is that these are the words we use for expressing assent and dissent. Under assent I include all of the following uses of the word "true." (They were originally distinguished by Professor P. F. Strawson in "*Truth*," *Analysis*, X, 1949, and in "*Truth*," *Proceedings of the Aristotelian Society*, Supplementary Volume XXIV, 1950. Professor Morris Weitz has summarized them in "*Oxford Philosophy*," *Philosophical Review*, LXII, 1953.) 1. The confirmatory use: A utters a value judgment and B says "That's true," confirming A's statement by adding the weight of his own opinion. 2. The admittive use: A utters a value judgment in a context in which B's statement "That's true" means "I admit it." Such a context would occur when A and B were disputing a point and A gave a convincing argument to B. 3. The concessive use: A utters a value judgment and B replies "That's true, but . . ." meaning "I concede what you say, but there are other things to be taken into account." 4. The agreeing use: A utters a value judgment, and B says "That's true," thereby indicating to A that he shares A's judgment. 5. The novelty use: A utters a value judgment which is surprising or new to B. B says "That's true," registering the fact that he had not thought of this before and that he agrees with A.

None of these instances of B's statement, "That's true," is an assertion that there is a correspondence between A's value judgment and an actual state of affairs in the world. What has led some philosophers to think that value judgments are true, and that consequently it makes good sense to talk about verifying them, is that we do use the word "true" in all these ways in reference to someone's uttering a value judgment. But in so using the word "true," we are merely expressing our assent. We are saying in different ways that we are in agreement with the value judgment that has been expressed. If the word "false" were used in similar contexts, we would be expressing our dissent or disagreement with what was being said. The words "true" and "false" are not being used here to claim that an assertion is verified or falsified, not even that it *could* be verified or falsified,

and still less that the method by which the assertion is to be verified is an empirical method.

What may be called the verification use of the word "true" has, as its criterion of application to any statement, the *intellectual or rational acceptability* of what is stated. And it is just in this sense that it is less confusing to speak of value judgments as being justified or unjustified than as being true or false. For there is a basic difference between what makes a value judgment intellectually or rationally acceptable and what makes an empirical assertion intellectually or rationally acceptable. The latter is intellectually acceptable (or true) if it can be empirically verified. The former is intellectually acceptable if there are good reasons for judging something to have a certain value. These good reasons do not consist in empirical evidence alone. A person can justify his decision or disposition to judge that something has a certain value only by a complex process of reasoning fundamentally different from empirical verification. If in establishing a value judgment as intellectually acceptable, we keep in mind how different this is from establishing the intellectual acceptability of nonevaluative assertions, then whether we call the value judgment "true" (in the verification use of the word, not merely in its assent use) is purely a verbal matter.

B. The process of verification

I come now to the central problem of this and the following three chapters. What sorts of reasons are good reasons in justifying value judgments? My purpose is to bring out the over-all logical structure of the justification of value judgments. Since I am here dealing with all types of judgments (moral, aesthetic, political, religious, matters of etiquette, and so on) and since I include both judgments according to standards and judgments according to rules, my discussion is bound to appear in some respects superficial. However, one specific point I wish to defend is that there are certain ways of reasoning which constitute a unified pattern of thought for justifying all value judgments. I am not interested in the particular points of difference that can be shown between, say, justifying moral judgments and justifying the judgments of art criticism. I shall indicate later what I think is the *locus* of such points of difference, but this

will only be done in order to throw light on the principles common to all justifications of value judgments.

I distinguish four general phases in the over-all process of justifying value judgments: verification, validation, vindication, and rational choice. We *verify* value judgments by appeal either to standards or to rules which we have adopted. We *validate* standards or rules (i.e., we justify our adopting certain standards or rules) by appeal to higher standards or rules. The adoption of standards or rules which themselves cannot be validated by appeal to any higher standards or rules results from our decision to accept a whole value system. We *vindicate* our accepting a whole value system by appeal to the way of life to which we are committed. Our commitment to a way of life can be justified in terms of a *rational choice* among different ways of life. It is because the three essential steps of verification, validation, and rational choice must follow the first step of verification that I have not been willing to use the term "verification" to cover the entire process of justifying a value judgment. To use the term "verification" for the first step, however, will not lead to difficulties if my foregoing remarks about verifying and justifying are kept in mind.

I adopt the terms "verification" and "validation" from Professor Kurt Baier's, *The Moral Point of View*. In Chapter 2 of that book, Professor Baier argues convincingly that, once we have adopted certain standards, our value judgments can be verified empirically. The statement that something fulfills or fails to fulfill a clearly defined standard is an empirical statement. The same thing is true of value judgments according to rules, since the statement that an act is in accordance with or violates a clearly defined rule (as distinct from the statement of the rule itself) is an empirical statement.

For both types of value judgment the *process of verification* is identical with the *process of evaluation* which I have described in Chapter 1. The only difference between verifying a value judgment and carrying out an evaluation in order to arrive at a value judgment lies in the circumstances in which the process takes place. The process itself remains the same. We speak of verifying a value judgment when someone has already arrived at a judgment and is asked to give reasons for it, or when the person himself becomes doubtful of his own judgment and wants to see whether, on reflection, he is

right in holding it. We speak of evaluating something when we are trying to come to a decision about its value, that is, when we are trying to arrive at a justified value judgment. In the first case we are trying to decide whether it is reasonable to accept what someone has offered as a correct opinion; in the second we are trying to form an opinion which it would be reasonable to offer as correct. In both cases we are trying to determine the correct opinion on a given matter (the value or disvalue of something). The process by which a judgment is verified is identical with the process of evaluation leading to that judgment because both are procedures of reasoning concerning the same judgment. To verify the judgment were to go through a procedure of reasoning which, if the judgment were not yet decided upon, would lead anyone validly to the judgment for its conclusion. To evaluate is to go through the same procedure for the purpose of arriving at a judgment which has not yet been decided upon.

The relation between the processes of evaluation and verification is like the relation between a student's proving a theorem and a teacher's correcting his proof. Professor Bernard Mayo's analysis of this situation is particularly instructive:

Let us suppose that the candidate is asked to prove a certain theorem of which he has not yet been given the proof, but he knows enough to work out the proof. Now it might be supposed that the reasoning is done by the candidate, who does the actual thinking out and writing down of the theorem; what the examiner does is merely to check the reasoning. But this would be a mistake. Reasoning is not the same thing as thinking or writing: it is the application of certain universal standards to the thinking and the writing; in this case, the principles of valid deduction. Obviously this is what the examiner does: he asks whether this proposition follows from that, whether there is a contradiction or inconsistency somewhere, whether what is claimed is a proof really is a proof. But this is not just what the examiner does. It is what the candidate does too. . . . The examiner judges someone else's performance; the candidate judges his own; and that is the only difference there is. That is why it would make sense for the examiner to say 'The right answer is the one which I should have given if I were doing the paper', and for the candidate to say 'The right answer is the one I should mark correct if I were examining the paper' and similarly for the wrong answer. (B. Mayo, *op. cit.*, p. 61.)

The relation between the evaluator and the verifier of a value judgment is parallel to this relation between candidate and examiner. The evaluator's purpose is to come to a conclusion about the value of something which can stand up under critical scrutiny. He is trying to arrive at a justifiable judgment. The verifier's purpose is to decide whether the judgment of the evaluator is justified. Both must appeal to the same rules of reasoning. These rules are the ordinary rules governing empirical verification.

When we verify a value judgment the process of reasoning is that outlined in Chapter I. It consists in finding out whether the evaluation does fulfill or fail to fulfill the *given* standards of evaluation to the degree stated in the judgment. Or, if the evaluation is according to rules, it consists in finding out whether the act, choice, or decision being evaluated is in fact required by, permitted by, or forbidden by the *given* rule. In the former case one must determine by empirical tests the good-making and bad-making characteristics of the evaluation (and, in the case of a ranking, the good-making and bad-making characteristics of the other members of the class of comparison). Then one must grade the evaluation according to the given standards (or one must rank the members of the class of comparison and decide where the evaluation falls in the resulting order of preference). In the case of evaluations according to rules one must determine by empirical tests whether the evaluation does come under the rule, and if it does, whether it complies with or violates the rule.

It is important to realize that, since no question is being raised here about the appropriateness of judging the evaluation by the given standards or rules, the *verifier's* pro-attitudes and con-attitudes need not be involved in the process of verification. Here is one point at which verification differs from evaluation. (This difference, be it noted, is not a difference in the rules of reasoning, which are still common to both processes.) Since the verifier *must* accept the given standards or rules in order to carry out his verification and since he is not concerned with the validation of such standards or rules, his own attitudes are irrelevant. But the evaluator has *chosen* to apply those particular standards or rules whose fulfillment is something he approves of, likes, or favors and whose nonfulfillment is something he disapproves of, dislikes, or disfavors. Whether his

pro-attitudes and con-attitudes are justified depends on whether the standards or rules are appropriate, that is, on whether they can be validated. But this is of no interest to the verifier.

C. The process of validation

If the foregoing were a complete account of the justification of value judgments, one could say that all value judgments are empirically verifiable and hence true or false in the same sense as non-evaluative assertions. But it is clear that we have not succeeded in justifying a value judgment merely by showing that the evaluatum does or does not fulfill certain standards or rules. Another question immediately arises. Are those standards or rules appropriate ones for judging an evaluatum of that sort? We must not only justify the claim that, given the standards or rules, the evaluatum has a certain value. We must also justify the application of those standards or rules in the given circumstances. This is where validation comes in. As Professor Baier puts it:

We have seen that value judgments can be verified just like factual claims, but that in value judgments we make claims that give rise to a further question, namely, whether the criteria employed are the right ones. Factual judgments are decisively confirmed if they are empirically verified. Value judgments, on the other hand, must be not only verified but also validated. It is not enough to show that, if certain criteria *are* employed, then a thing must be said to have a certain degree of 'goodness'; we must also show that these criteria *ought* to be employed. . . . A remark ceases to be a value judgment and turns into a factual claim as soon as the question of the *appropriateness* of the *criteria* is rejected as unnecessary or irrelevant. (K. Baier, *op. cit.*, pp. 75-76.)

What does the process of validation consist in? By way of a preface, let us notice that the standard or rule which is being validated is universal; it must apply to a given class of things. We cannot say that an object *X* is good according to a standard *S* unless all objects which are similar to *X* in the relevant respects are good. (The relevant respects are those features of an object in virtue of which it fulfills *S*, in other words, its good-making characteristics.) Similarly we cannot say that doing an act *A* in circumstances *C* is right according to a rule *R* unless all acts of type *A* done in circum-

stances *C* are right. The class of all objects which are good (or bad) according to a standard *S* defines the *scope* or *range* of *application* of *S*. The class of all acts which are right (or wrong) according to a rule *R* likewise defines the *scope* or *range* of *application* of *R*. In accordance with what was said in Chapter I, the scope or range of application of *R* is the class of all acts "falling under" or "covered by" *R*.

1. THE DEMAND FOR VALIDATION

In order to understand the various ways in which validation can be carried out, let us first see how the demand for validation arises. I shall consider in turn the validation of a standard and the validation of a rule. Suppose someone claims that a painting he is looking at is a good painting because (among other things) its colors are harmonious. The following dialogue represents schematically how the demand for validation of a standard may arise:

- A: This is a good painting.
 B: Why?
 A: Because, for one thing, its colors are harmonious.
 B: Why does that make the painting a good one?
 A: Because having harmonious colors is a good reason for judging this to be a good painting.
 B: Why is it a good reason for judging this to be a good painting?
 A: Because having harmonious colors is a good reason for judging paintings of kind *K* to be good, and this painting is of kind *K*.
 B: I admit that this painting is of kind *K* and that its colors are harmonious, but why is having harmonious colors a good reason for judging paintings of kind *K* to be good?

B's last question may be stated in another way—why is having harmonious colors a valid or appropriate standard for judging paintings of kind *K*? To answer this question satisfactorily, *A* must show not only that he, but that anyone (and therefore *B*), is justified in adopting the given standard for judging paintings of kind *K*. That is, he must give good reasons why paintings of kind *K* are correctly included in the range of application of the standard.

This goes beyond the verification of *A*'s value judgment "This is a good painting." The demand for verification only requires *A* to show

that this painting fulfills the standard of having harmonious colors. When we verify (or falsify) the value judgment, we *assume* that the standard of having harmonious colors correctly applies to the K class of paintings. This assumption is precisely what is in question when we demand a complete justification of A's value judgment. In order to justify the judgment completely it is necessary to justify this assumption. And this requires that we validate the standard.

It should be noted that the same argument holds for those aestheticians and art critics who claim that each work of art is unique and can be judged only in terms of standards peculiar to it. If the appropriateness of these standards is challenged, the evaluator must give reasons to justify his appealing to them rather than to other standards (unless he is willing to say that his appeal to them is purely a matter of personal taste which he cannot defend as objectively valid). To say that a certain standard is peculiar to one object is to say that it is not appropriate or valid to use that standard when judging any other object. The class of things to which the standard validly applies is a "unit class" (a class having only one member). Now the claim that a given standard applies uniquely to one object is still a claim that must be justified. How can that be done? I submit that there are only two possible methods a person can use, if he is to be fully rational in justifying his claim. He can appeal to a higher standard (or rule) in terms of which he can *validate* the standard in question as applicable to only one object. Or (the more likely method in this situation) he can *vindicate* his use of the standard in terms of his whole way of life. In the latter case he must show that the adoption of the given standard is part of a way of life in which works of art are appreciated in a certain way. The standard would then be shown to be an outward expression of personal aesthetic taste, that is, it would be shown to be the standard a person *would* apply to a work of art, and to that work of art alone, if he had a certain aesthetic taste. A *full* justification of the standard would require, as we shall see, that the whole way of life which includes that aesthetic taste be shown to be a *rational* way of life.

I should like to add, however, that I do not think the claim that *all* aesthetic standards apply uniquely to single works of art can finally be defended. Surely there are standards that apply, say, to (the class of) representational paintings but do not apply to (the

class of) nonobjective or abstractionist paintings. One such standard could be that we see nature in a new light, or learn to look at things in a new way, once we come to understand fully how the artist himself has "represented" nature in his painting. In the case of pure abstractions, this standard would not be appropriate. But there are standards that do not apply to such broad classes of comparison as all representational paintings or all abstractionist paintings. There are standards that may apply only to representational paintings of a certain period. Thus we may apply a certain standard only to works of the High Renaissance; for example, the degree to which the artist presents to us an idealistic vision of human dignity or nobility within a Christian framework. There are standards that may apply only to the works of one artist—for example, how developed are the "painterly" qualities of Titian's works, or how free from sentimentality are the "mystical" aspects of Tintoretto's works. That is to say, we might make our class of comparison *only* the works of Titian or *only* the works of Tintoretto and apply standards that we claim to be appropriate to one of these classes of comparison and not to the other. (Standards of freedom from sentimentality might themselves differ when we judge Tintoretto and when we judge, say, Renoirs.) Furthermore, if an artist has painted different versions of one subject matter (e.g., the Annunciation) we might use standards in judging these paintings that apply validly only to the particular subject matter in question. In such arts as music, drama, and the dance, we can narrow the range of application of a standard still further, without yet reaching a "unit class" of application. We might use standards which we hold to be appropriate for judging the various performances of one work. We might apply certain standards to, say, Handel's oratorio, "The Messiah," and then judge all the known performances of that work by appeal to those standards. Finally we *may* use certain standards as applicable to a single performance of one work (or to a single painting, sculpture, poem, or building). In this last case alone would a standard be considered uniquely applicable to a particular object. Whether there are such standards is a question requiring detailed investigation in aesthetics and so beyond the scope of my present discussion. I only maintain here that not *all* aesthetic standards are of this sort.

The demand for validation of a *rule* may be represented as follows:

- A: This act is wrong.
 B: Why?
 A: Because it violates rule *R*.
 B: Why is the fact that it violates rule *R* a reason for condemning it?
 A: Because rule *R* forbids any act of kind *K*, and this act is of kind *K*.
 B: I admit that rule *R* forbids acts of kind *K*, that this act is of kind *K*, and therefore that rule *R* forbids this act. But why is that a reason for condemning the act?
 A: Because rule *R* is a valid or appropriate rule for judging acts of kind *K*. (Or: Because acts of kind *K* are correctly included in the range of application of rule *R*.)
 B: Why?

This last question demands that *A* validate rule *R*. *A* is asked to give good reasons for anyone's adopting *R* as a rule covering acts of kind *K*. If the rule is validated, then it is shown to be a reason for *condemning* acts of kind *K*, which violate it. Similarly it is shown to be a reason for praising (admiring, liking, or having any other pro-attitude toward) acts of a kind which are covered by the rule and which fulfill it. (There is a parallel in the case of standards. Validating a standard is justifying anyone's adopting it as a basis for evaluating objects of a certain kind. Hence the standard becomes a basis for taking a pro-attitude or con-attitude toward a given object of that kind.)

As in the case of standards, giving a complete justification of a value judgment made according to a rule must go beyond the verification of it. In order to verify it one must *assume* that the rule correctly includes the class of comparison in its range of application. One can then verify the judgment by seeing whether the act in question (the evaluatum) complies with or violates the rule. In order to justify the judgment completely, however, one must justify making this assumption; one must *validate* the rule.

2. THE THREE STEPS OF VALIDATION

How is a standard or rule to be validated? Its complete validation requires three logical steps. 1. We must show (by methods to be discussed below) that the standard or rule is *relevant*. That is,

its scope or range of application must include the class of comparison of the given value judgment. As a consequence of this first step, the evaluatum is shown to belong to a *class* of things which are correctly judged by the standard or rule. 2. We must show that neither the circumstances in which the evaluatum occurs nor anything out of the ordinary about the evaluatum (i.e., anything which distinguishes the evaluatum from the other members of the class of comparison) permits us to make an *exception* to the general application of the standard or rule determined in step 1. As a consequence of this second step, the evaluatum itself, not merely a class of things to which it belongs, is shown to be correctly judged by the standard or rule. 3. We must show either (a) that no other valid standard or rule conflicts with the one being applied, or (b) that, if there is a conflict, the one being applied takes precedence over all those in conflict with it.

If we fail in any of these three tasks, our validation of the standard or rule is not complete. When we successfully accomplish step 1 there is a certain presumption that the standard or rule is valid for judging the evaluatum. When we successfully accomplish step 2 there is a greater presumption that the standard or rule is valid. But we have conclusive reasons for believing that the standard or rule is valid only when we have completed all three tasks. Only then do we *know* that the standard or rule is valid.

If someone should ask on what grounds such a claim to knowledge can be made, my answer would be that no other way of reasoning could yield better reasons for accepting a standard or rule as valid. I am merely trying to explicate the pattern of reasoning which would yield the *best* results for justifying a value judgment, taking into consideration only the nature of evaluation and the *point* of demanding such a justification in everyday life. I am not trying to set up a special method which will lend support to a particular theory of value (such as intuitionism or naturalism), nor am I selecting a method which will make a certain predetermined set of value judgments turn out to be true. *Whatever* value judgments are considered to be true by *anyone* (in any culture or in any period of history), good reasons for accepting the standards or rules on which they are based must be put forward. What better reasons could there be for accepting a standard or rule than the successful

accomplishment of the three steps of validation I have outlined? If someone is not satisfied that such reasons would logically be the best reasons, he must show that an alternative method would constitute a better way of validating a standard or rule. I submit that this cannot be done, if I am right in my foregoing analysis of evaluation and value judgments. (For further consideration of this question, see Chapter 4, Section B.)

I have not yet made clear, however, in what way the three steps of validation can be carried out, and this must be done before the reasonableness of the entire process of validation can be fully demonstrated. By what methods can the three tasks listed above be accomplished? There are three possible methods. The first and second are alternative methods for carrying out step 1, and the third must be used for carrying out steps 2 and 3. Let us designate the standard to be validated as *S* and the rule to be validated as *R*. The three methods may then be summarized as follows:

Method I: Appeal to standards or rules which are more general than *S* or *R* and from which *S* or *R* can be deduced.

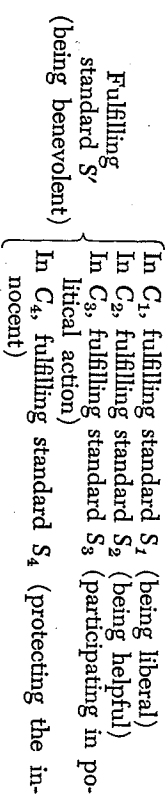
Method II: Appeal to standards for judging the consequences of fulfilling or of not fulfilling *S* or *R*.

Method III: Appeal to standards or rules for deciding whether it is better to make an exception to *S* or *R* than to follow *S* or *R*, and for deciding the relative precedence of any standards or rules which are in conflict with *S* or *R*.

I shall now consider how the first two methods can be used to accomplish the task of step 1 and how the third method can be used to accomplish the tasks of steps 2 and 3.

STEP 1. One way to show that *S* is relevant—that it is correct to include a given class of comparison and hence a particular evaluation in the range of application of *S*—is to appeal to a more general standard *S'* from which *S* can be deduced. (The relevance of *S'* would itself have eventually to be validated by Methods I or II.) When do we say that one standard can be deduced from a second? The statement that an object fulfills the second standard in certain circumstances must logically entail the statement that that object fulfills the first standard in those circumstances (though not necessarily in *all* circumstances). An example or two will clarify this.

Consider the relation between the standard of benevolence and the standard of liberality in giving to charity. A person fulfills the standard of benevolence when, in certain circumstances, he is liberal in giving to charity. But he also fulfills the standard of benevolence when, in other circumstances, he helps a person in distress (as in an automobile accident), and when, in a third set of circumstances, he joins a movement for racial equality, and when, in a fourth set of circumstances, he goes out of his way to speak up in behalf of a man being unjustly defamed. This may be put in the form of a simple diagram:



If we adopt standard *S'* as validly applicable to an object *X* (say, any human being), then the statement "X fulfills *S'*" entails the statements "X fulfills *S*₁ in *C*₁"; "X fulfills *S*₂ in *C*₂"; and so on. When this relationship between *S'* and *S*₁, *S*₂, . . . *S*_n holds, I shall say that *S*₁, *S*₂, . . . *S*_n are *deducible* from *S'*. Of course, if *S*₁ is deducible from *S'*, then the fact that an object *X* fails to fulfill *S*₁ in *C*₁ entails its failure to fulfill *S'*. However, the failure of *X* to fulfill *S*₁ in circumstances *other than* *C*₁ does not entail such a failure. In a given set of circumstances to be benevolent requires one to be liberal in giving to charity; in a different set of circumstances (say, when the individual himself is destitute), it does not. In these circumstances, failure to be liberal does not mean failure to be benevolent.

It is to be noted that the deducibility of *S*₁, *S*₂, . . . *S*_n from *S'* depends on the adoption of *S'* as a *valid* standard for judging a class of objects *K* of which *X* is an instance. If the validity of *S'* were brought into question a new process of validation would be necessary. In the given case perhaps the standard of benevolence as applied to all human beings (class *K*) could itself be deduced from a more general standard, such as the principle of brotherly love. Thus we might have the following logical pattern:

Fulfilling S'' (brotherly love)	$\left\{ \begin{array}{l} \text{Fulfilling } S_1' \text{ (being benevolent)} \\ \text{Fulfilling } S_2' \text{ (being honest)} \\ \text{Fulfilling } S_3' \text{ (being just)} \\ \text{Fulfilling } S_4' \text{ (being conscientious)} \end{array} \right\}$	$\left\{ \begin{array}{l} \text{In } C_1, \text{ fulfilling } S_1 \\ \text{In } C_2, \text{ fulfilling } S_2 \\ \vdots \\ \text{In } C_n, \text{ fulfilling } S_n \end{array} \right\}$
---	---	---

Here the logical relation between S'' and S_1', S_2', \dots, S_n' would be the same as that between S_1' and S_1, S_2, \dots, S_n . The applicability of S'' would have to be assumed, just as the applicability of S' was assumed before. If this assumption were questioned, a still higher stage of validation would have to be carried out, and so on, until we reached the supreme norms of the value system we have adopted. (As we shall see later, any further justification would require that we shift from validation to vindication.) To sum up, if a man is judged to be morally good because he gives liberally to charity, one can justify the relevance of the standard used, by showing that it follows from the more general standard of benevolence as a sign of good character. The claim that benevolence is a standard relevant to a moral judgment of good character is itself in need of justification; if it is justified, then it follows that being liberal in giving to charity is also validated as a relevant standard.

Let us take another example. We can show that having harmonious colors is a valid standard for judging paintings of a certain kind (say, European paintings of the seventeenth and eighteenth centuries), by deducing it from a more general standard which is appropriate to paintings of that kind. Such a general standard might be the "integration" or "organic unity" of all the formal and qualitative elements in the composition of a painting. If we accept this standard as relevant, and if color harmony is one of the conditions necessary for a certain kind of painting to fulfill the standard, then we have established that color harmony is a relevant standard.

The relevance or applicability of a standard may also be established by Method II. One might argue, for instance, that liberality in giving to charity is a relevant standard for judging a man's character, because the over-all effects of people's acts are *better* when it is fulfilled than when it is not. In other words, the more widely the standard is adopted in a society, and the more completely it is

fulfilled by those who adopt it, the *more ideal* the society will be. (Method II would not normally be used to validate aesthetic standards.) The judgment that the effects are better or that the society is more ideal is itself a value judgment made according to some *other* standard. This other standard may be the maximizing of happiness for as many people as possible, or the prevention of social conflict, or the attainment of economic security, or the like. Whatever it is, the original standard in question has not been shown relevant by Method II until this other standard has been established (by Method I or II). However, assuming that this standard is relevant (for judging the social consequences of adopting and fulfilling other standards), the relevance of the original standard is established.

Methods I and II may also be employed to validate a rule. The methods must show that a given act is legitimately included in the range of application of the rule. Suppose we judge an act to be wrong because it violates the rule: always pay the bills you have incurred as a result of your purchases. On what grounds can we show that it is correct to apply this rule to a person when he does not pay his bills? It is not sufficient merely to say that he had a charge account at the store and charged things when he purchased them. This is to *assume* the relevance of the rule that one must pay for things one has charged. Method I may be used to establish the relevance of the rule as follows. The rule in question is a particular case of a more general rule (that of keeping one's promises). To incur a debt contextually implies that one promises to repay it—unless this promise were made, there would be no debt. To be in debt *means* to be obligated by one's implicit or explicit promise to repay it. Now when one opens a charge account at a store, by that very act one promises to pay one's bills in the future. (This is part of what the social practice of opening a charge account means. One does not have to *say* that he promises to pay his bills.) By Method I it is correct to condemn a person's act on the ground that it violates a rule if that rule can be deduced from a more general rule which covers the act. In the present example the given act violates the rule: one must pay one's bills. There is a more general rule which covers the act in question and from which the aforementioned rule can be deduced, namely: one must keep one's prom-

ises. Therefore it is correct to condemn the act as a violation of the more particular rule.

To argue the case by Method II, we would imagine what it would be like if people in general did not repay their debts, and then to predict (infer) what consequences would probably result from this practice. We then judge these consequences to be bad according to a certain standard, such as the standard of maximizing happiness. Again, we would assume that this standard is itself relevant to the case. If this assumption were challenged, the validation of the standard would have to be carried out by using Method I or II.

The use of Methods I and II in validating a rule reveals how a social practice as a whole can be justified. A social practice is a way of behaving according to a set of rules. These rules define the practice, so that to follow the rules is to engage in the practice. "It is the mark of a practice," he says, "that being taught how to engage in it involves being instructed in the rules which define it, and that appeal is made to those rules to correct the behavior of those engaged in it." (J. Rawls, "Two Concepts of Rules," *The Philosophical Review*, LXIV, 1, 1955, p. 24.) In these cases, an act of a certain sort would not even be described as an act of that sort unless it fell under the rules which define a practice. A person could not be said to be voting in a political election, for example, unless he placed a proper kind of mark on an official ballot at the appropriate time (between certain specified hours on Election Day) in the authorized place (the district election center), unless he was registered in a recognized party, unless the person he voted for was a candidate running for a definite office and had been duly nominated for that office, and so on. To describe an act of checking a piece of paper as an act of voting, we must understand the act as part of a practice defined by a particular set of rules. "To engage in a practice, to perform those actions specified by a practice, means to follow the appropriate rules. If one wants to do an action which a certain practice specifies then there is no way to do it except to follow the rules which define it." (Rawls, *op. cit.*, p. 26.)

It is important to notice that one might violate the rules and still be engaged in the practice. Indeed, if the act is correctly described

as a violation of the rules, one *must* be engaged in the practice. To be engaged in the practice only means that one's acts fall under the rules and are describable in terms of the rules, not that they are always in accordance with the rules. In order to say that a person keeps a promise (or repays a debt), we have to presuppose that the person is engaged in the practice of making promises (or incurring debts). The same holds true when we say that he breaks a promise (or does not repay his debts). He must still be engaged in the practice of making promises (or incurring debts), otherwise his acts do not fall under the rules defining the practice.

Now the fact that the rules are needed to describe the practice does not of itself imply that they are validated or that the acts are *correctly* judged according to them. To engage in a practice is to have one's acts fall under certain rules and consequently to make it possible to judge one's acts by the rules. But having one's acts fall under the rules still needs to be justified, and justifying it requires that the rules be shown to be relevant in judging the acts. To show that the rules are relevant is to *justify the entire social practice*. Let us see why this is so.

When an act is judged according to rules which define a practice, the judgment is verified (but not validated) when the practice is specified and the act is correctly described as one that complies with or violates a rule of the practice. The judgment "That act is wrong" is verified when the practice of promising, for instance, is specified, and when the act is correctly described as breaking a promise. This does not fully justify the judgment, however; one might ask why the practice of promising should be engaged in at all. This is a demand for the justification of a social practice. In order to justify engaging in it (and so acting under and being judged by certain rules) it is necessary to validate the rules which define it. That is, it is necessary to justify *adopting* those rules in the first place and so *making* one's acts fall under them. When Methods I and II are used to validate practice-defining rules, then Methods I and II are used to justify social practices. To validate such a rule by Method I as that one must repay his debts, we must show that the social practice defined (in part) by that rule is included in a wider or more general social practice (say, making promises) which is itself justified. A social practice *P* is included

in a wider or more general social practice P' when every act which is describable in terms of P is also describable in terms of P' , but not all acts describable in terms of P' are describable in terms of P . (That is, to engage in P is always to engage in P' also, but not vice versa.)

Use of Method I in justifying a social practice is most familiar to us in the field of ethics, where it constitutes the basic idea of restricted utilitarianism. We validate the rules defining a certain social practice by means of Method II when we give a moral justification for the practice by appealing to the principle of utility (or the greatest happiness principle. In such a case the rules defining the practice are themselves evaluated. They are judged as having instrumental value or disvalue to an inherently valuable end (the greatest happiness of all sentient beings). In Chapter I I mentioned the case of judging rules themselves with regard to their stringency or importance, according to their instrumental value. Validating practice-defining rules by Method II is an instance of this.

Professor Stephen Toulmin's, *An Examination of the Place of Reason in Ethics* (Cambridge, England: Cambridge University Press, 1950), contains an excellent account of this distinction between justifying a particular act falling under a social practice and justifying a social practice as a whole. Professor Toulmin gives the following example. I have borrowed a book from someone and I am asked why I ought to return it.

... I may reply that I ought to take it back to him, 'because I promised to let him have it back before midday'—so classifying my position as one of type S_1 . 'But ought you *really*?' you may repeat. If you do, I can relate S_1 to a more general S_2 , explaining, 'I ought to, because I promised to let him have it back.' And if you continue to ask, 'But why ought you *really*?' I can answer, in succession, 'Because I ought to do whatever I promise him to do' (S_3), 'Because I ought to do whatever I promise anyone to do' (S_4), and 'Because anyone ought to do whatever he promises anyone else that he will do' or 'Because it was a promise' (S_5). (S. Toulmin, *op. cit.*, p. 146.)

Professor Toulmin then adds that beyond this point we cannot raise the question, 'Why ought I (you) to do this particular act in these circumstances?' because 'there is no more general 'reason' to be

given beyond one which relates the action in question to an accepted social practice." (*Ibid.*, p. 146.)

Suppose, however, that the circumstances are changed, so that I seem to be justified in making an exception to the social practice. Then a second sort of reason must be introduced. I cannot merely appeal to the social practice to justify my action. Professor Toulmin gives this example:

... If I have a critically ill relative in the house, who cannot be left, the issue is complicated. The situation is not sufficiently unambiguous for reasoning from the practice of promise-keeping to be conclusive. I may therefore argue, 'That's all very well in the ordinary way, but not when I've got my grandmother to look after: whoever heard of risking someone else's life just to return a borrowed book?' (*Ibid.*, p. 147.)

There is yet a third kind of reason in ethics. One can reply to a demand for justification, not of an exception to a social practice, but of the social practice itself. Suppose I know that, *given* the practice of promise-keeping, I am obligated to return the book; I want to know, however, why such a practice *ought* to be accepted. I am no longer asking for justification of a particular act of promise-keeping. I am asking for a justification of promise-keeping itself.

To question the rightness of a particular action is one thing: to question the justice of a practice *as a practice* is another. (*Ibid.*, p. 149.)

How would such a question arise? Professor Toulmin contends that it would arise under changing economic, social, political, or psychological conditions.

If a society has a developing moral code, changes in the economic, social, political or psychological situation may lead people to regard the existing practices as unnecessarily restrictive, or as dangerously lax. If this happens, they may come to ask, for instance, 'Is it right that women should be debarred from smoking in public?', or 'Would it not be better if there were no mixed bathing after dark?', in each case questioning the practice concerned *as a whole*. (*Ibid.*, p. 149.)

The kind of reason which would then become appropriate, according to Professor Toulmin, is the kind of reason I have placed under Method II. It consists in using a standard for making a value judg-

ment of the consequences of the social practices involved. Here is where the principle of restricted utilitarianism occurs in ethics.

The answer to be given will . . . be reached by estimating the probable consequences

(i) of retaining the present practice, and
(ii) of adopting the suggested alternative.

If, as a matter of fact, there is good reason to suppose that the sole consequences of making the proposed change would be to avoid some existing distresses, then, as a matter of ethics, there is certainly a good reason for making the change. (*Ibid.*, pp. 149-150.)

This distinction between moral reasons for doing a particular act and moral reasons for adopting a social practice is by no means a recent discovery in ethics. Professor Toulmin himself points out that Socrates, in the *Crito*, was making this very point.

It was this distinction between the 'reasons' for an individual action and the 'reasons' for a social practice which Socrates made as he waited for the hemlock: he was ready to die rather than repudiate it—refusing, when given the chance, to escape from the prison and so avoid execution. As an Athenian citizen, he saw that it was his duty (regardless of the actual consequences in his particular case) to respect the verdict and sentence of the court. To have escaped would have been to ignore this duty. By doing so, he would not merely have questioned the justice of the verdict in his case: he would have renounced the Athenian constitution and moral code as a whole. This he was not prepared to do. (*Ibid.*, p. 151.)

The conclusion I wish to draw from Professor Toulmin's discussion is this. When Method II is used to validate rules which define a social practice, validating the rules is the same as justifying the social practice as a whole. This may readily be distinguished from justifying a particular act falling under the social practice. We justify a particular act by verifying a value judgment of the act according to the practice-defining rules. When we shift from justification to validation in justifying that judgment, we shift from justifying an act falling under a social practice to justifying the social practice itself. (It will be shown in Chapter 5 that social practices may also be justified by a process of vindication.)

Throughout this account of validating a standard or rule by

Methods I and II, I have been making two assumptions. First, that the evaluatum of the value judgment which is being justified does not constitute an exception to the standard or rule. And second, that there are no other standards or rules which conflict with the one being validated, or that if there are, the one being validated takes precedence over them. A consideration of the second and third steps of validation brings these assumptions into question.

STEP 2. An act may be a case of breaking a promise, but that is not a *conclusive* reason for our judging it to be wrong. This is so even though we realize that the agent himself must be engaged in the practice of making promises in order to be said to break a promise, and even though we believe that the practice of making promises is a justified practice. The given act might be a legitimate exception to the rule that we must keep our promises. Although the rule *generally* covers acts of that kind, the *given* act may be distinctive in a relevant respect. An example of this is to be found in a passage I have quoted from Professor Toulmin. A person is excused from returning a borrowed book because he must look after his critically ill grandmother. If keeping his promise to return the book involves risking his grandmother's life, we say it is right for him to make an exception to the rule: Always keep your promises. The range of application of the rule ordinarily includes promise-breaking acts. But because of the circumstances, this particular promise-breaking act becomes permissible, even obligatory. The judgment that the act is wrong cannot be justified; the evaluatum, in those circumstances, does not fall under the rule according to which the judgment is made.

How do we know when the violation of a rule (or the nonfulfillment of a standard) can be counted as a legitimate exception? On what grounds do we justify making such an exception? The answer is to be found in the use of Method III. We must appeal to a standard or rule which permits such an exception to be made. *Assuming* the validity of *this* standard or rule, it is possible to decide the question. A full justification of the resulting value judgment, however, would necessitate the justification of that assumption. We would have to validate (or vindicate) the standard or rule appealed to in making the exception. Validation of this standard or rule would be made by further use of Methods I, II, and III.

It is easy to see how a standard can function as a ground for making an exception to another standard (*S*) or to a rule (*R*). If, according to the standard, it is *better* to make the exception than to fulfill *S* or *R*, then making the exception is justified. We are dealing here with a value ranking whose class of comparison is composed of 1 anything that fulfills *S* or *R* and 2 something which, in the given circumstances, would constitute a failure to fulfill *S* or a violation of *R*. When the result of evaluating the members of this class of comparison is such that the evaluatum in 2 is judged to have more value or less disvalue than instances of 1, then it is legitimate to make an exception to *S* or *R*. Let us consider an example of justifying in this way an exception to *S*, and then an exception to *R*.

Suppose a new library building is being planned for a city. One of the architectural principles (standards) accepted by the city officials is that of functionalism—the best building is the one that most effectively serves the purposes for which it is to be used. According to this standard, good-making characteristics of a library would include such things as spaciousness of stacks and reading rooms, quietness, efficiently organized offices, easy availability of books, and so on. Let us suppose, furthermore, that the best library building as judged by the standard of functionalism would cost so much to build that the city would have to take funds allotted to another project (say, slum clearance) to pay for it. The choice becomes: fulfill the standard of functionalism and damage the slum clearance project, or fail to fulfill the standard and preserve the project. An evaluation is then made of these two alternatives *according to a standard*, such as the welfare of the people of the city. An evaluation according to this standard might result in ranking the second alternative as better than the first. An exception would then be made to the standard of functionalism in architecture. If an architect were to present two plans for the library building, one of which was clearly better than the other according to the standard of functionalism, the better one would *not* be chosen. It would not be considered really better in the given circumstances. Hence an exception to a standard is justified on the basis of the circumstances in which the evaluatum occurs. The standard of what is “really better” is the standard of the public welfare. If this standard

were itself to be challenged, the process of validating it by Methods I, II, and III would have to take place. (Or else it would have to be vindicated.) In this example I am assuming that this standard is not challenged.

The most familiar case of making an exception to a rule on the basis of a standard occurs when we judge that the *effects* of following the rule will be worse in a given set of circumstances than the effects of violating it. Thus we sometimes make exception to a moral rule, such as the rule forbidding lying, on grounds of a moral standard, such as that of minimizing suffering. Exceptions to rules of etiquette or of custom are quite frequently made on such grounds. In all instances of this sort an evaluation is being made which results in the (ranking) judgment that violating a rule is better than following it. Whenever this occurs, making an exception to the rule is justified (assuming, of course, the validity of the standard according to which the evaluation is made).

How can a *rule* be used to justify an exception to a standard (*S*) or to another rule (*R*)? The pattern of reasoning is essentially the same as in the preceding cases. We carry out an evaluation to see whether it is better not to fulfill *S* or to violate *R* than to fulfill *S* or comply with *R*. But here the evaluation is made according to a rule, not a standard. Justifying an exception to a standard on the basis of a rule might occur as follows. Suppose a businessman has set an ideal standard of efficiency for his company. He usually judges to be good whatever fulfills this ideal, bad whatever fails to fulfill it. Sometimes in order to fulfill his ideal most completely he might well have to break a law (e.g., the minimum wage law). From the point of view of his ideal, it would be better to break the law than to obey it. In these circumstances, however, he makes an exception to his standard. If he is to follow the rule which requires him to obey the law, he must do certain things which work against his ideal. Assuming the rule to be valid, he is justified in making such exceptions to his standard.

Making an exception to a rule *R* is justified on grounds of another rule whenever the other rule takes precedence over *R* and requires an act which is forbidden by *R* (or forbids an act enjoined or permitted by *R*). For example, when we think that obeying a certain law is a serious violation of our moral code, we may believe that we

are justified in breaking the law. We do not have to condemn the law *as a law* on moral grounds, as a conscientious objector would condemn a military draft law as a law. We might generally approve of the law and yet in unusual circumstances believe it is right to break it. The moral rule "Always help a person in distress" might require our breaking a traffic law (say, going through a red light in order to get a critically injured person to a hospital).

It is worth repeating that in all these ways of justifying exceptions to standards and rules, some other standard or rule is presupposed. One can always challenge a person's claim that he is justified in making an exception by demanding that the presupposed standard or rule itself be justified. Such a justification process may consist either in validating the standard or rule by further use of Methods I, II, and III, or in vindicating its adoption. How it can be vindicated will be discussed in Chapter 5.

Having shown how an exception can be justified, I have also shown what it means to say that making an exception is *not* justified. Making an exception to *S* or *R* is not justified when doing so has less value or greater disvalue than fulfilling *S* or *R*. Making the exception must be ranked in relation to not making the exception. Such a ranking must appeal to some standard or rule other than *S* or *R*, which in turn must be either validated or vindicated.

STEP 3. So far we have analyzed the first two steps of the process of validation. In step 1, we show that in ordinary circumstances a given evaluatum is correctly included in the range of application of a standard *S* or a rule *R*. In step 2, we show that the evaluatum does not occur in such unusual circumstances, or differ in such a way from other members of the class of comparison, that it constitutes a legitimate exception to *S* or *R*. Now a third step is necessary if the validation of *S* or *R* is to be complete; for there might be other standards and rules which are in conflict with *S* or *R* and which take precedence over them.

What does it mean to say that standards and rules may be in conflict? Two standards conflict with each other when something which is good according to one is bad according to the other. Two rules are in conflict when an act which is in accordance with one violates the other. There is a conflict between a rule and a standard when an act which fulfills the standard violates the rule, or when

an act which is in accordance with the rule fails to fulfill the standard.

In order to resolve a conflict between standards or rules logically, one must decide which standard or rule takes precedence over the other. There are two ways in which this can be done: by appeal to a second-order standard and by appeal to a second-order rule. Let us examine first the appeal to a second-order standard *S'*. A first-order standard *S*₁ takes precedence over another first-order standard *S*₂ if and only if something which completely fulfills *S*₁ is *better* (according to *S'*) than something which completely fulfills *S*₂. Similarly a rule *R*₁ takes precedence over another rule *R*₂ if and only if an act done in accordance with *R*₁ is better (according to *S'*) than an act done in accordance with *R*₂. In the case of a conflict between a rule *R* and a standard *S*, *R* takes precedence over *S* when an act done in accordance with *R* is better (according to *S'*) than an act which completely fulfills *S*. If an act which completely fulfills *S* is better than an act done in accordance with *R*, then *S* takes precedence over *R*. In each instance we determine whether one thing is *better* than another by ranking the two things according to a second-order standard (*S'*).

There is a parallel situation when the relative precedence of standards and rules is decided by appeal to a second-order *rule* rather than a second-order standard. In this case, the act in accordance with the standard or rule taking precedence is said to be *right*, while the act in accordance with the other standard or rule is *wrong*. The rightness and wrongness of the act are determined by an evaluation according to a second-order rule *R'*.

It should be noted that when one rule *R*₁ is found to take precedence over another rule *R*₂ (whether by appeal to a second-order standard or by appeal to a second-order rule), we can say that *R*₁ makes a "higher claim" upon us, lays down a "heavier obligation," or imposes a "more stringent duty" than *R*₂. (Which of these phrases is most appropriate depends on whether the type of rules concerned are moral, legal, etiquette, and so on.) We might also say it is better that our conduct be guided or regulated by *R*₁ than by *R*₂. Similarly when a standard *S*₁ is shown to take precedence over another standard *S*₂, we can say it is better for someone to try to fulfill *S*₁ than *S*₂. For the sake of simplicity I have given definitions only of

conflicts between two standards, between two rules, and between a standard and a rule. Conflicts may occur, however, among a great number of standards and rules. In these cases all the standards and rules can theoretically be arranged in a hierarchy, each one (except the lowest) taking precedence over another. Whether the actual set of standards and rules accepted by an individual ever constitutes such a neat "value system" is a matter I shall deal with shortly.

The standard by which we judge whether something which fulfills S_1 or R_1 is better than something which fulfills S_2 or R_2 (and accordingly whether S_1 or R_1 takes precedence) I have called a second-order standard. By a second-order rule, we judge that an act which fulfills S_1 or R_1 is right and an act which fulfills S_2 or R_2 is wrong (and accordingly we determine that S_1 or R_1 takes precedence over S_2 or R_2). A standard or rule is second-order simply in virtue of the fact that it is used for judging the relative precedence of other (first-order) standards or rules. Such second-order standards and rules function exactly like first-order standards and rules in a process of evaluation. For any second-order evaluation, the members of the class of comparison are *ideals* of S_1 and S_2 (i.e., things which perfectly fulfill S_1 and S_2) or ideals of R_1 and R_2 (i.e., acts which are in perfect accordance with R_1 and R_2). If the conflict holds between S_1 and S_2 , then the class of comparison of the second-order ranking has two members: the ideal of S_1 and the ideal of S_2 ; if the conflict is between S_1 and R_1 , the class of comparison is the ideal of S_1 and the ideal of R_1 ; and so on. When many standards and rules are in conflict, the class of comparison is composed of the ideals of all these standards and rules.

When a second-order evaluation is made according to a standard, its purpose is to arrange the ideals of all the first-order standards and rules in an order of preference depending on their good-making and bad-making characteristics. The first-order standards and rules that are in conflict do not determine the good-making and bad-making characteristics. For according to each of them there is one member of the class of comparison which is *ideal* in comparison with which all the others are in varying degrees bad. It is the second-order standard that determines the good-making and bad-making characteristics on the basis of which some of these ideals are judged to be better than others. When a second-order evalua-

tion is made according to a rule, the class of comparison is always composed of acts, decisions, or choices that have been judged as ideal according to the first-order standards or rules (whose relative precedence is being decided). The second-order evaluation is a process of judging the rightness or wrongness of each act, decision, or choice in the class of comparison by appeal to the second-order rule.

A person's second-order standards and rules reflect his general outlook on life. If we know that a man lives according to the second-order rule "Whenever there is a conflict between moral rules and the rules of long range self-interest (prudence), do what is prudent, not what is moral," we know something fundamental about the person. His second-order rule reveals to a great extent his "philosophy of life." If a man's basic goal in life is to become wealthy, he will judge the various ideals of his first-order standards and rules according to their instrumental value in realizing that goal. He will subordinate such ideals as obeying the law, enjoying a happy married life, and being well liked to the more fundamental ideal of gaining wealth. The latter standard takes precedence over all other standards and rules in his life.

Two further points should be noted about second-order standards and rules. In the first place, they are the logical, not necessarily the psychological, basis on which an individual organizes his life. They are the standards and rules which an individual *would* specify as his basic principles or "values" *if* he were thoroughly honest and *if* he reflected objectively about his own life. In the second place, it is only the unusual person who integrates all his standards and rules in an order of relative precedence. Most people do not have such organized views of life. At any one time they might have several second-order standards or rules which themselves conflict; according to one, S_1 takes precedence over S_2 , and according to another S_2 takes precedence over S_1 . These second-order standards or rules might also shift from time to time, so that a person gives up some life goals and adopts others as his situation in life changes. It is perfectly possible, indeed, for a person to live entirely without second-order standards or rules. In that case when his first-order standards and rules come into conflict with each other (as they are bound to do in the everyday affairs of life) he will act in an incon-

sistent manner. At one time he will try to fulfill one standard, at another time a conflicting one. Today he will follow rule R_1 at the cost of violating R_2 ; tomorrow he will (in similar circumstances) follow R_2 at the cost of violating R_1 .

The clearest examples of second-order standards and rules are to be found in the lives of those "integrated" persons who have dedicated themselves to a fundamental ideal. I have mentioned the case of the man who seeks wealth above all things. Other examples might include a deeply committed religious man, a political fanatic, a thoroughgoing militarist, a professional athlete in training, a creative musical genius, and a dedicated scholar or scientist. At the core of such individuals' lives are certain integrating second-order standards and rules according to which at least some, if not all, of their first-order standards and rules are arranged into a hierarchy. If such a person were asked to justify his (implicit or explicit) belief that one first-order standard or rule takes precedence over another, he would cite certain second-order standards (e.g., standards of total commitment in advancing a political cause, or standards of impartiality and logical rigor in the pursuit of truth) and certain second-order rules (e.g., "I must always obey the will of God, or I must never do anything detrimental to my development as a tennis player").

Are there third-order standards and rules by which to resolve conflicts among second-order standards and rules? Theoretically, there are. They would be appealed to if a validation of a second-order standard or rule were demanded and if that validation included the use of Method III. Probably very few people actually organize their lives on the basis of such third-order standards or rules, however.

So far I have discussed the use of second-order standards and rules in the life of an individual. They also have an important place in the functioning of a social group, whether that functioning be organized (as in the case of an army waging a battle, a company producing, distributing, and selling goods, or a university carrying on its academic functions) or relatively unorganized (as in the case of a cocktail party or a race riot). Sometimes the functioning of a group is so organized as to constitute a social practice. The first-order rules which define a social practice may be arranged hier-

archically according to second-order rules (which are then as much a part of the definition of the practice as are the first-order rules). When the first-order rules of a practice conflict with each other, it is necessary to decide their relative precedence. If a social group is to be able to function at all, there must be an established way of resolving conflicts among its rules and this requires an appeal to second-order rules. Suppose an army unit has the two rules: "Always obey the commands of a superior officer," and "No one is to enter this building without a special pass." If an officer without a special pass commands the soldier on duty to let him enter the building, which rule ought the soldier obey? A well-run army unit will specify the relative precedence of such conflicting rules by means of a second-order rule. After the statement of the rule requiring a special pass, there might be added the second-order rule, "Absolutely no personnel of any rank will be exempted from this requirement."

The third step in validation, to conclude, is accomplished by Method III. Here the task is to show that the standard or rule to be validated is not in conflict with any other validated standards or rules, or if it is, that it takes precedence over them. The relative precedence of conflicting standards and rules is found by ranking them according to second-order standards or rules. These second-order standards or rules themselves must be assumed to be validated (or vindicated) if the original value judgment is to be justified. Should this assumption be challenged, the second-order standards or rules would have to be validated by Methods I, II, and III (or they would have to be vindicated by the method to be discussed in Chapter 5).

D. On the argument for radical skepticism

If we disregard for the moment the possibility of vindicating standards and rules and consider the justification of value judgments only in terms of verification and validation, then the following argument arises. In order to justify a value judgment we must not only verify it by appeal to standards and rules but we must also validate those standards and rules. Since such validation itself depends on the justifiability of higher standards and rules, then the latter must also be validated. This would require our accepting still higher

standards and rules, whose validation in turn would presuppose still higher standards and rules, and so on. Now either this process goes on to infinity or it stops somewhere. In the former case, since the ultimate grounds on which a value judgment is based can never be reached, no value judgment can ever really be shown to be justified. In the latter case, the point at which validation stops must be arbitrarily chosen, since it cannot itself be validated. The conclusion of radical skepticism is then drawn: we cannot legitimately claim to *know* what is good or bad, right or wrong.

The force of this argument derives from viewing evaluative reasoning on the model of mathematical deduction. Value judgments are thought of as theorems that are strictly entailed by statements of standards and rules. The latter in turn are thought of as theorems strictly entailed by statements of higher standards and rules, until we arrive at axioms, that is, statements of *ultimate* standards and rules. Since these axioms are not themselves entailed by any higher statements, they are held to be pure postulates, arbitrarily chosen. For this reason all value judgments are claimed to be equally absurd.

This argument does not claim that value judgments are matters of taste, about which there is no disputing. In matters of taste we do not even *attempt* to argue for or against a person's statements. If value judgments were considered to be matters of taste, no claim would be made that a value judgment is justified or that a standard or rule has rational grounds. As soon as the question "Why should I accept the standards and rules you appeal to?" is raised, we are no longer concerned with a matter of taste. The radical skeptic is not denying that we raise such questions about value judgments. He grants that they can be as rationally grounded as mathematical theorems. But, he says, *finally* they are absurd because they rest on axioms, which are decision-statements, not claims-to-truth. As a result, they are beyond the boundaries of rational discourse. One "value system" (a set of standards and rules arranged hierarchically according to their relative precedence) is no more justified than another.

An attempt is made to preclude this kind of skepticism by those who claim that one set of axioms is intuitively certain. Such a claim rests, I think, on the following line of reasoning. Since the axioms

are statements of standards or rules, they are not tautologies—they are not true by definition or by the law of noncontradiction. On the other hand they are not empirical statements. They are normative, not descriptive. Now because they are not tautologies, they must be synthetic statements. And because they are not empirical, they are not true a posteriori but must be true a priori. Therefore they are synthetic a priori statements. How then do we know whether they are true? Since axioms by definition are not deducible from any higher premises, they are not known to be true because other statements are known to be true. And we have just seen that they are not known to be true on the basis of evidence provided by our senses. They must therefore be known to be true by an immediate intuition, yielding a unique sort of knowledge that is totally independent of all other knowledge. It is knowledge of what ought to be, not knowledge of what is. Consequently value judgments *can* be justified. Their justification consists in three things—verification by reference to standards or rules, validation of these standards or rules by reference to higher standards and rules (finally by reference to ultimate or highest standards and rules), and synthetic a priori intuition of the truth of these ultimate standards and rules. Accordingly we do have knowledge of what is good or bad, right or wrong. There is nothing arbitrary, irrational, or absurd about our value judgments.

I shall be brief in my criticism of these arguments. Their basic error is to misconstrue the nature of our reasoning about value judgments. They do this for two reasons. They consider verification and validation purely in terms of reasoning in the empirical sciences and in mathematics, respectively. And they do not take into account the two further stages of justifying value judgments: vindication and rational choice.

Both the argument for and the argument against radical skepticism rest on the false assumption that there are only two ways of reasoning about anything—the way we reason in mathematics and the way we reason in the empirical sciences. Both arguments accept the fact that, unlike empirical assertions, value judgments must be validated as well as verified. But it is assumed that the only alternative to empirical verification is deductive inference from axioms. Since validation is different from empirical verification, it must con-

sist in a final appeal to axioms. As a result, two questions arise which both the skeptic and his opponent are trying to answer. What are the axioms from which value judgments are deduced and on which they ultimately depend for their justification? Are these axioms such that we can claim knowledge on behalf of our value judgments?

The analysis of validation which I have given in this chapter lends some support to this way of looking at the problem. I have emphasized how the validation of any standard or rule depends on our accepting a higher standard or rule. And it is perfectly true that, in this relation of a given standard or rule to a higher one, the first is deducible from the second (in just the ways I have been describing under Methods I, II, and III). It would seem, then, that if we are not to go on to infinity we must arrive at standards or rules that are ultimate. What is there left but to say that these ultimate standards or rules are either arbitrary (and hence are accepted without reason) or are knowable only by a synthetic a priori intuition?

A fundamental mistake is being made here. Certainly there is a deductive element in justifying value judgments, just as there are deductive elements in establishing scientific theories and laws. But this does not mean that evaluative reasoning is like mathematical reasoning in *all* respects. We are overlooking what is *distinctive* about evaluative reasoning. We are trying to fit it into some preconceived logical pattern, such as the analytic-synthetic dichotomy. If we look at the way we actually do reason about what is good or bad, right or wrong, we find that there is an important difference. Just as such reasoning does not stop at verification of a value judgment but goes on to the validation of standards and rules, so it does not stop at validation. There are two further steps which are *essential* in the complete justification of any value judgment. These are the vindication of whole value systems and the rational choice of a way of life. The next three chapters are concerned with these steps in the justification of value judgments.

Value systems and points of view

4

A. The concept of a point of view

In order to understand the third and fourth steps in the justification of value judgments, namely vindication and rational choice, it is first necessary to distinguish two concepts and see how they are related to each other. These are the concept of a value system and that of a point of view. A value system is a set of standards and rules of *a certain kind* arranged according to the place they have in the verification and validation of value judgments (and prescriptions) of *that kind*. Thus a *moral* value system is a set of moral standards and moral rules that are appealed to in verifying moral judgments and moral prescriptions, and that are arranged in an order of relative precedence corresponding to the hierarchy implicit in their validation. An *aesthetic* value system is a set of aesthetic standards and aesthetic rules that are appealed to and validated in the justification of aesthetic judgments and aesthetic prescriptions. In like manner there are political value systems, economic value systems, religious value systems, etiquette value systems, prudential value systems, and so on. In each case, the structure of the system is logically determined by the justification of value judgments and