

*The Study of Human Nature and the
Subjectivity of Value*

BARRY STROUD

THE TANNER LECTURES ON HUMAN VALUES

Delivered at
The University of Buenos Aires

June 7, 1988

BARRY STROUD is Professor of Philosophy at the University of California, Berkeley. He was born in Toronto, Canada, and attended East York Collegiate, the University of Toronto, and Harvard University. He has spent research leaves in Oxford and in Venice and has held visiting academic appointments in Great Britain, Canada, Norway, Mexico, Australia, and the United States. In 1986–87 he was John Locke Lecturer at the University of Oxford and Visiting Fellow of New College and of All Souls College. He is the author of *Hume* (1977) and *The Significance of Philosophical Scepticism* (1984) as well as numerous essays on a variety of philosophical topics. He is a fellow of the American Academy of Arts and Sciences.

I believe that the philosophical study of moral and other values is filled with difficulty. Our best attempts to understand evaluative thought of all kinds —all questions of good and bad, better and worse —seem to me to distort or threaten to obliterate the very phenomenon we want to understand. But if we refrain from pressing for a philosophical expose of values we appear to ourselves to be simply acquiescing in a way of thinking and acting without understanding it. And that leaves us dissatisfied. So we persist, and end up misrepresenting and so still not understanding the phenomenon of value.

I would like to present at least the outlines of the dilemma I see. Having it clearly before us is a necessary step toward finding some way out of it. In these two lectures I can explain it only sketchily and at a regrettably high level of generality.

I stress that it is a philosophical way of thinking about values in general that I am interested in, not any particular morality or political arrangement or set of values in itself. It is a very powerful conception of what is really going on when human beings deliberate, evaluate courses of action, and make choices, or assess the choices and actions of others. If some such conception really is at work in our understanding of ourselves, it can be expected to affect the way people think concretely about what they are doing, and why. And it can come in that way to affect what people actually do. There are perhaps good reasons in general to doubt that such an abstract, purely philosophical theory could ever have such palpable effects. But on the other hand it seems hard to deny that many of the ways we think and speak about our current social arrangements, and the justification typically offered for them, do rest on some such conception of value in general. I will not have time to go into the question of the extent to which that is really so.

I will first try to describe the main outlines of the conception of evaluative thought that I have in mind and then identify the kind of distortion or denial I think it leads to. Then, in the second lecture, I will turn to the question of how and why we are so inevitably driven toward that dead end. I think it comes from nothing more than our desire to understand ourselves in a certain way.

I

Any attempt in philosophy to understand morality or evaluative thought generally leads almost inevitably to what I shall call “subjectivism.” It is not always easy to notice this tendency, let alone to lament it, since the kind of view one is led to appears to have gained the status of orthodoxy. There seems to be no other way to think about values. And so, we think, nothing true is being distorted or denied at all.

The idea, in a word, is that values are “subjective,” that questions of value are not questions with “objective” answers, that the goodness or badness of a thing or a course of action is not something that belongs to the world as it is in itself, independently of us. There are many different versions of this single thought. I will not be concerned with each one of them, It is what they all have in common that leads to the difficulties I see.

What they all have in common is the thought that there are no evaluative facts. In general, when we say or believe something, if things are the way we think they are, if the world is in fact the way we say it is, then what we say or believe is true. When what we say is false, things are not that way, the world is not in fact as we say it is. In science and all other forms of inquiry we seek the truth. By that I mean nothing lofty, abstract, or metaphysical. I mean only that in this or that particular way we want to find out what is so, how things are, what the world is like in one or another respect. The question can be quite particular and trivial (e.g., Where is that book I was reading yesterday?) or extremely general and profound (e.g., What, if any, are the fundamental ele-

ments of the universe?). It is always a question of what is so — what are the facts. Whether or not something is so, whether there is anything there to be discovered or not, is in general something that holds quite independently of whatever we might happen to think about it, however we might feel about it, or even whether we are at all interested in it or not.

On the “subjectivist” view, matters of value — of the goodness or badness, the beauty or worth, of a thing or action — are not in this way anything to be found among the facts of the world at all. They are therefore not part of anything that scientific or any other kind of cognitive investigation could study and try to make progress on. There is in that sense no possibility of moral or aesthetic or, in general, evaluative knowledge. Not because our faculties are too weak to discern the true value of things, and not even because evaluative matters are so complex that we can never expect universal or even widespread agreement about them. It is, rather, that in the realm of values there is simply no “objective” truth to be known. The world in itself is just what it is; it is simply there. It is the totality of facts, and it is value free.

Of course, human beings do take an interest in certain facts. They care about certain things and not about others, they want certain things, they try to bring about certain states of affairs and to prevent others. Those are undeniably facts of the world. Human beings are part of the world, and they do think and feel and act in those ways. In short, human beings value some things or states of affairs more highly than others. That is a fact of the world, but it is not an evaluative fact. It is not a matter of one thing’s being better than another. It is simply a matter of human beings’ regarding one thing as better than another. For the “subjectivist,” there are “objective” facts of what humans do, but not of the value of what they do, or of the value of anything else.

What “subjectivism” denies, then, is not that human beings do place value on certain things, but only that there is such a thing as being correct or incorrect in those valuing, as we can be correct or

incorrect in our beliefs about the facts. When we say or think that something is good or worthwhile, or evil or ugly, our thinking or saying it is certainly something that is so, but either there is nothing at all that makes what we think true or false, or if in some way there is, it is only something about us, something “subjective.” And it is nothing evaluative.

Hume put the view this way:

Take any action allow'd to be vicious: Wilful murder, for instance. Examine it in all its lights, and see if you can find that matter of fact, or real existence, which you call *vice*. In which-ever way you take it, you find only certain passions, motives, volitions and thoughts. There is no other matter of fact in the case. The vice entirely escapes you, as long as you consider the object. You never can find it, till you turn your reflexion into your own breast, and find a sentiment of disapprobation, which arises in you, towards this action. . . . So that when you pronounce any action or character to be vicious, you mean nothing, but that from the constitution of your nature you have a feeling or sentiment of blame from the contemplation of it. Vice and virtue, therefore, may be compar'd to sounds, colours, heat and cold, which, according to modern philosophy, are not qualities in objects, but perceptions in the mind.¹

Hume thought this “discovery” about the nature of morality was a great advance in the study of human nature. It was the “discovery” (although Hume, of course, was not the first to make it) or what I am calling the “subjectivity” of value.

Hume thought that not only values and colors are “subjective,” but also, most famously, causality itself. Given the way human minds work, we will inevitably come to believe in necessary causal connections between some of the things we experience. But nothing in the world corresponds to that belief. “Necessity is

¹ David Hume, *A Treatise of Human Nature*, ed. L. A. Selby-Bigge (Oxford: Clarendon Press, 1958), pp. 468–69.

something, that exists in the mind, not in objects,” he said.² We think causal necessity is something “objective,” but it is really nothing but a “subjective” “determination of the mind.”³ This famous treatment of the idea of causality can still serve today as our best model of “subjectivism.” Other more recent varieties can all be measured against it.

There are many different positive versions of the “subjectivist” idea. For Hume, in speaking of necessity there is really nothing to speak of except what is in your own mind. And, as he says, when you “pronounce” upon the value of something “you mean nothing, but that . . . you have [a certain] feeling or sentiment” toward it. Taken literally, that implies that value judgments are really just statements of feeling. That particular idea is not essential to “subjectivism.” Another version says that if you say that something is vicious you are not stating that you have got a certain feeling but rather are simply expressing or giving vent to a feeling or attitude you have toward the thing. Your remark is like a cheer or a sigh and is therefore neither true nor false.⁴ Or you might be both stating facts about the action and expressing a feeling toward it. Another version says that you are reporting or expressing a feeling and also encouraging others to have that same feeling or attitude.⁵ For some “subjectivists” feelings are not involved at all; in “pronouncing” upon the value of something you are recommending or prescribing it, not saying anything that is true or false of it.⁶ A quite different kind of theory holds that when you say that something is vicious you are saying only that the thing is such as to produce certain feelings or experiences or desires in human

² Ibid., p. 165.

³ Ibid.

⁴ See, e.g., A. J. Ayer, *Language, Truth, and Logic* (New York: Dover Publications, 1946), chap. 6.

⁵ See, e.g., C. L. Stevenson, *Ethics and Language* (New Haven: Yale University Press, 1958), chaps. 2, 4, 9.

⁶ See, e.g., R. M. Hare, *The Language of Morals* (Oxford: Clarendon Press, 1952), chaps. 1, 12.

beings of such-and-such kinds. Whether the thing does or would have such effects is a straightforward matter of fact. But for the “subjectivist” there is nothing evaluative in the facts such judgments state. There couldn’t be. They speak only of nonevaluative effects to be brought about in human beings by certain “objective,” nonevaluative states of affairs.

“Subjectivism” carries with it a certain view of moral discussion or disagreement. It cannot see it as a dispute as to how things are, or what is so. Those who dispute about whether it is better to do X or to do Y when it is not possible to do both do not dispute about any matter of fact. Of course, they might disagree about certain facts as well, but the purely evaluative dispute is not factual. The disputants’ valuing or attitudes or feelings are opposed to each other, so that at most only one of them can prevail, but the one who does prevail cannot be said to be getting things right while the other is getting them wrong. The one who prevails gets, or gets more of, what he values. But their dispute, if it is evaluative, is not a dispute about whether the world is such that X is better than Y or that Y is better than X.

The theory obviously has great appeal. It is extremely widely believed, in one form or another. In fact, it can seem to be the only kind of account there could be, largely because it alone among all theories avoids what would otherwise be an apparently insoluble problem. If values were part of the “objective” world, what sort of thing could a value be? How could there be such a thing as an evaluative fact or state of affairs? We know that where a thing is, what shape it is, how much it weighs, even what color it is or how much it costs, even whether human beings want it or get pleasure from it — are all matters of “objective” fact. But how could there be an additional fact to the effect that the thing is good or bad, or better than something else? The unintelligibility or “queerness” of what values would apparently have to be if they were “objective” has been one of the strongest arguments for “sub-

jectivism.” As befits a metaphysical theory, it is defended on what are really metaphysical grounds.

The theory also has its moral or political appeal. It seems to express something to which we attach positive value —the idea that nothing or nobody can push me around in matters of evaluation. There is no position from which one person’s values can be criticized as incorrect or misguided. Nor is a person’s choice of what to do or the best way to live constrained by some “objective” standard against which it can be measured. The thought that the world cannot force us to accept one set of values rather than another can be liberating. It does not necessarily make life easy. There are great differences and conflicts among people’s valuing, and social and political life is a matter of resolving those conflicts and reconciling opposed interests. But what calls for solution is the question of which is to prevail. Each opposing interest must somehow be accommodated. All are there to be dealt with, and there are none that can be dismissed on the grounds that they are mistaken.

I have called what is common to all forms of “subjectivism” a metaphysical theory. It involves a conception of what the world is really like —a specific, determinate idea of the nature of “objective” reality. It is a world that lacks some of the things that most people appear to believe it contains. It is in that sense a more restricted world than what we seem to accept in everyday life. For Hume it contained no necessary causal connections between events, and no colors or sounds. No causal sentences or color sentences were true of the world. For the “subjectivist” about values no evaluative sentences are true of the world, even though we appear to say and think that some things are good, or are better than others. Evaluative thoughts or beliefs or attitudes are part of the world, but there is nothing in the world that makes those thoughts true or false. All such evaluative facts have been eliminated from the subjectivist conception of what the world is like.

Eliminating something from our conception of the world is in ordinary circumstances a familiar procedure. We could be said to be doing it whenever we find out that something we used to believe is not so. This happens every day in small matters and, over longer periods of time, cosmically. Great scientific breakthroughs are sometimes needed to bring about an altered conception of the world. Other, smaller changes take less. But in every case those old ways of thinking are then abandoned.

With the metaphysical theory of “subjectivism” things are different. Human beings—even “subjectivists”—continue to talk about and appear to believe in those very things that the theory claims are not really part of the world. There is a sense in which they are not abandoned. We cannot help getting experiences of color and believing on that basis that objects around us are colored. We do inevitably come to value certain things more than others. The “subjectivist” philosopher of human nature says that those things we inevitably perceive and come to believe in are not in fact to be found in the “objective” world. But any such theorist, being human, will inevitably get those very perceptions and beliefs that the theory says are only fictions and cannot be true. The “subjectivist” will inevitably believe that grass is green, for example, while also holding that no object in the world has any color. And he or she will regard a particular murder as vicious or bad while also insisting that no value statements are true, that the viciousness or badness of something is nothing in the world.

This seems to require of “subjectivism” both detachment from and engagement with the very same experiences, ideas, and beliefs. We must stand apart from our color beliefs and our evaluations while also holding onto them. Given the force with which the world inevitably operates on us, this would seem to make reflection on the austere, restricted reality of “subjectivism” at best unstable—a momentary grasp of what you take to be the way things really are, from which your humanity immediately rescues you, plunging you back into a rich world of colors and vice and virtue

which reflection had apparently revealed to be nothing but illusions generated only by your own constitution. No one has given more poignant expression to this plight, while remaining in the grip of both sides of it, than Hume.⁷

But it is not just a matter of psychological instability, or oscillation. It is a question of whether that restricted view of the fully “objective” world can even be reached. That is the question I want to ask. Can we coherently think of a world in which all our valuings are exposed as only “subjective”? Could we then continue to understand ourselves to be making any evaluations at all? I think neither defenders nor opponents of “subjectivism” have taken this question seriously enough.

We say how we think the world is by saying what we believe to be so. But as long as we simply specify how things are, or how we take them to be, we will never arrive at the view that I am calling “subjectivism.” In fact, if we tried to specify all the things we believe, and we took that list to express our conception of what the world is like, what we believed would be incompatible with “subjectivism.” One thing I believe is that grass is green; another is that some acts are vicious murders, that the deliberate killing of a human being is a very bad thing. If I take these beliefs to express part of my conception of the world, I will have to conclude that it is a fact, or part of the way the world is, that grass is green, that some acts are vicious murders, and that the deliberate killing of a human being is a very bad thing. My conception of the world will not then be “subjectivistic” about colors or values. So at the very least the “subjectivist” account of the world must not include the contents of any of those beliefs. In saying how things really are it must not mention the colors of things or their value.

But merely leaving such things out of one’s conception of the world is not enough in itself to express the “subjectivist” concep-

⁷ See especially *Treatise*, book 1, part 4, section vii.

tion. To leave certain features out of my conception of the world is not necessarily to conceive of a world which lacks those features. I might concentrate for some reason on only certain aspects of things. For example, I might think only about the size of the objects in my house, without mentioning their location, where I got them, or how much they cost. But that does not mean that I think that only their size is real, that they do not really have any location, any origin, or any cost. Similarly, I might specify a huge number of physical facts about the movements of particles, the presence of certain forces in the world, and so on, without mentioning the colors of anything. But I do not thereby imply that I think things have no color. I simply say nothing about their color one way or the other. And if I say only that certain physical movements occurred and the effect was the death of a human being, I say nothing about the value of what went on, but I do not imply that it was not in fact a vicious murder, or that I believe it was not. So merely stating some of one's beliefs about the world without mentioning the colors of things, or their value, does not automatically make one a "subjectivist" about colors or about values. Leaving something out is not the same as saying that there is no such thing.

"Subjectivism" clearly needs the thought, then, that colors, or values, or whatever is said to be purely "subjective," are *not* part of the world. Rather than merely conceiving of a world without conceiving of colors or values, it must conceive of a world which lacks colors and values. It involves a claim of exclusiveness. The negative claim about what the world does *not* contain is as essential to "subjectivism" as the positive claim about what the world is really like.

But "subjectivism" also requires the thought that people nevertheless do have beliefs about, or experiences of, those very features which it holds are not part of reality. The point of calling the source of those beliefs or experiences merely "subjective" is that we only *think* things are that way, or we have experiences which

we wrongly *take* to represent the way things are. Without that, there would be nothing to be a “subjectivist” about. The theory is a theory about human thoughts or beliefs or experiences. So it cannot deny that we have such thoughts and beliefs and experiences.

If all this is what “subjectivism” requires, how is it to be shown that “subjectivism” is true in a particular domain? How is it to be shown, for example, that there is nothing in the world corresponding to our beliefs about colors or about values—nothing to make them true or false? With a theory like Hume’s it can look easy. He thought that all that was available to us in perception of the world were momentary, independent atoms of sensory information. Anything we ever think about must somehow be constructible out of such meager data. The task of his science of man was to explain how we develop our elaborate conception of the world with so little information to go on. Given only such restricted data, various features of our own minds will obviously have to play a large role. To the extent that our own mental operations alone can explain the origin of ways of thinking that go beyond what is available in the minimal data, those ways of thinking will be seen to have a wholly “subjective” source. The world would not have to contain anything corresponding to those ways of thinking in order for them to arise quite naturally in us as they do.

This is a strategy that many “subjectivist” philosophers since Hume have made use of, and continue to make use of today. If you can explain how people come to think or experience something without having to suppose that those thoughts or experiences represent anything that is so in the world that gives rise to them, you will have exposed the thoughts or experiences as “fictions” with a wholly or partly “subjective” source. “Objective” reality would therefore include no more than what is found to be essential for explaining everything that happens, including human beings’ getting the thoughts and beliefs and experiences we know they get. So to say that colors are not part of the world,

or that nothing in the world as it is in itself has any value, would be to say that nothing like “Grass is green” or “That was a vicious murder” has to be taken to be true in order to explain why people come to think that grass is green or that an act was a vicious murder. No colors would need to be ascribed to anything in the world in order to explain people’s color perceptions and beliefs. And no values would need to be ascribed to anything in the world in order to explain why human beings value things as they do.

I call such explanations, if they are successful, “unmasking” explanations. They unmask or expose some of our beliefs or experiences as illusory in the sense of not actually representing the way things are in the world, even though it is perfectly understandable why we inevitably get such beliefs or experiences, given what we are like and the way the world works on us. Whatever we cannot help regarding as true in order to explain our thinking and experiencing what we do must be reckoned as part of the way the world is. Those indispensable beliefs about the world will not then have been exposed or unmasked by an explanation of their origins. On the contrary, they will have been vindicated. They will have been shown to represent things as they really are. But for all the rest, the world is not really the way they represent it as being.

This might be called an explanatory test or criterion of reality. The world as it is in itself amounts to all, but only all, those truths that are sufficient to explain what is so. Anything that is not needed for that explanatory purpose is not to be reckoned as part of the way things are.

I have said that this is one possible route to “subjectivism” about values, or about colors. It seems to rely on a certain faith in the simplicity of the universe. It sees the world as highly efficient and economical, as no richer than it needs to be for the explanatory purposes of science. I do not want to speculate about the origins of such a faith. Nor will I go into the details of any particular attempts to establish “subjectivism” by an appeal to

unmasking explanations. I can only say that it seems to me extremely implausible to think that they alone could do the job. They seem to work best, as in Hume's case, when you have already arrived at a restricted conception of what the world really contains. But establishing "subjectivism" in a particular area is a matter of arriving at that appropriately restricted conception of the world in the first place.

I want to turn away from all questions about how "subjectivism" about values might be established and look instead at what must be an essential ingredient in any form of the view, however it is arrived at. There must be some way of understanding the presence of what those unmasking explanations, if they were appealed to, would be supposed to explain or unmask. The "subjectivist" view of the world, for all its zeal in eliminating certain features we unreflectively seem to think are there, still must acknowledge as part of the world all those perceptions, beliefs, and attitudes of human beings which it claims have only a "subjective" source. And there is a question of how the presence of those perceptions, beliefs, and attitudes is to be understood.

The question arises as much for colors as it does for values, and it will be helpful to look at that case first. To entertain the view that colors do not belong to the "objective" world, but are at best projected onto, or falsely believed or perceived to be present in, a world that does not really contain them, we must ourselves attribute no color to anything (since we say there is none in the world) while nevertheless believing that there are many perceptions of and beliefs about the colors of things. The question is whether we can do that. It obviously depends on what perceiving colors or believing that things are colored amounts to, and on what it takes for us to understand that such psychological phenomena occur. If we are "subjectivists," it cannot depend on our supposing that any of the contents of such perceptions or beliefs are actually true of the world they are about. Can we make sense of the perceptions or beliefs if we no longer make that assumption?

We can, of course, understand people to have many beliefs about, and perhaps even perceptions of, things which we ourselves know do not exist. I understand that people believe in, and sometimes describe themselves as seeing, for example, ghosts or angels. People also think about centaurs and golden mountains. There is no doubt that such psychological phenomena occur. The explanation traditionally offered for our understanding of such facts relied on a simple compositional theory of thought. The concept of a ghost or a centaur was said to be a complex idea. It represents nothing that exists in the world, but it is a compound made up of simple elements, some of which do indeed find counterparts in the world. Our attributing thoughts or beliefs about non-existent things to others therefore does not require that we ourselves believe the world to be populated with the things those complex ideas represent. We can see how people come to think that way without our agreeing that the thoughts they have are true.

Even this theory does not completely sever our understanding of the thoughts of others from all our own beliefs about the way things are. The presence in our common world of objects like horses' bodies or the heads of men (or other even simpler things) is what enables us to think about such things and to attribute thoughts with those contents to the minds of our fellow humans. But even if that theory is perfectly satisfactory for thoughts about centaurs or golden mountains — which I do not believe it is — it would be of little help in explaining how the “subjectivist” can understand the presence of thoughts and perceptions of color. Surely our idea of color cannot be built up out of simpler elements' that are not themselves colors at all. Perhaps some particular colors or shades can be understood as mixtures or combinations of other colors or shades, but there are no “elements” which are not colors but which somehow could be combined in thought or experience to give us the idea of color in the first place.

Particular shades of color have traditionally been thought to be so simple that we can all understand what it is to have a per-

ception of them simply by having such perceptions. It has been suggested that we each understand in our own case what it is to have a perception of green, say, simply by perceiving green; we know what that is like. If we could each understand it in that way, we could perhaps then say that what others have when they have a perception of green is just the same as what we have. We know in that way what green is, so we know that the feature that others perceive when they perceive green is that same feature. And that is what we all ascribe to objects when we believe that they are green. This seems to involve no ascription of green to anything in the world and yet to acknowledge the presence of perceptions of green and beliefs about green things on the part of human beings constituted more or less as we are.

I do not find this traditional theory plausible, for reasons I will only state and not develop. I believe we could never come to understand in that first-person way what it is to have a perception of green. The theory says that having a perception of green, or perhaps several of them, is enough to teach us what having a perception of green is. But simply having perceptions of green could never be enough. There is no way of being directly acquainted with something, or simply gazing at or experiencing a particular item, and from the mere occurrence or presence of the thing somehow coming to understand it as a thing of a certain sort rather than of some other sort. That is what we must do if we are to understand something as a perception of green rather than, say, as a remarkable event, which it might also be. Nor is there any possibility, on the sole basis of "having" it, of understanding that we have got the same sort of thing this time as we have had before. Every two things are the same in some respect or other, and also different in countless respects, so whether we have got the same kind of thing on a second occasion depends on which respects are relevant and which not. And that cannot be fixed by an original item about which we understand nothing but which we merely "have."

Some surroundings are needed to make a thought into the thought of a certain kind of thing, so some surroundings are needed to make a thought about a psychological occurrence into a thought about a perception of green, say, and not something else. But if in trying to supply the surroundings needed to ascribe perceptions of particular colors to perceivers we find that we ourselves must also ascribe colors to some of the things we take them to be perceiving, we will have abandoned the “subjectivist” conception of reality. We will be conceiving of the world as containing colored things. The “subjectivist” thought must leave room in the world for perceptions of and beliefs about color, but the price of our understanding such things to be part of the world would be our also taking the world to contain colored things. Color perceptions and beliefs could not then be unmasked as illusory or as having no counterparts in the way things are. So “subjectivism” could not be established.

It might seem that that is not so, since there is at least one version of “subjectivism” on which it remains true that objects are colored. It says that what is ascribed to objects when we apply color words to them is a disposition to produce perceptions of color in appropriately placed perceivers in certain specified conditions. Objects really do (or do not) have such dispositions. So on that view our beliefs about the colors of things would indeed describe things as they are in the world. We would not be precluded from truly ascribing colors to objects.

But that dispositional theory does not really avoid the difficulty. It explains what it is for an object to be colored in terms of perceptions of color, but it says nothing about what a perception of color is, or what it takes for us to understand that there are such things as perceptions of color. Nor does it explain what a perception of green in particular is a perception of. The greenness involved in perceptions of green — what makes them perceptions of green — cannot itself be equated with a disposition to

produce perceptions of green, even if the greenness of objects can be explained that way.

That dispositional account of the greenness of an object makes essential use of an idea of green that cannot in turn be explained in that same dispositional way. So it must hold that there are perceptions of green even though no objects in the world possess that feature that they are perceptions of. And that is the same problem that faced other versions of “subjectivism” about colors. It must explain how we can understand particular perceptions to be perceptions of, say, green and not something else, while at the same time we hold that no objects in the world possess that feature that they are perceptions of.

When we think about what actually happens in everyday life, it seems that we constantly do rely on the public accessibility of such states of affairs as the greenness of grass in ascribing perceptions of greenness to our fellow human beings. We attribute color to objects in the world as a condition of attributing particular contents to perceptions. If that is so, and inescapably so, we will not be able to think of the world in the way “subjectivism” requires.

We can now see, I hope, a parallel difficulty for the “subjectivity” of values. This is where the threat of distortion or denial comes in. Those who think that a particular act was an act of murder and was vicious or wrong seem to have a certain thought about that act: they think it was wrong. Perhaps they think in general that the deliberate killing of a human being is a very bad thing. We can speak of such persons as having certain moral views or beliefs or opinions (in this case not very controversial). “Subjectivism” cannot deny that people have such views. It must insist on the fact. The question is how it can acknowledge and understand that fact while also holding that no such thoughts are ever true or false of the world.

I have said that there are many different positive theories of evaluative judgment which are all compatible with the negative

“subjectivist” thesis that values are not part of the “objective” world. One view is that the assertion of “That act was wrong” reports the presence of a certain feeling that the speaker has toward that act (for Hume, a “sentiment of disapprobation”). Another view holds that what is being said is that the act is such that all human beings of certain kinds would get a certain feeling toward it if they knew of it. Both these views see the so-called evaluative judgment as a factual assertion about actual or hypothetical feelings on the part of certain human beings. In that respect they are like the dispositional analysis of an object’s color.

One merit of all theories of this sort is that they preserve one striking feature of our evaluative thought. They allow that our reactions to the world do involve genuine beliefs about the goodness or badness of things. They see us as asserting what we take to be truths about the world. And there is very good reason for insisting that we think of our moral judgments as either true or false. Not only do we seem to believe them and assert them and try to support them by reasoning. Moral sentences can also be embedded in other sentences in what certainly looks like a purely truth-functional way.

For example, from the sentence “That act was wrong” and the sentence “If that act was wrong then whoever did it deserves to be punished” it follows logically that whoever did that act deserves to be punished. Any view which says that moral or other evaluative judgments are not assertions or are not, strictly speaking, true or false has great difficulty in accounting for that logical implication. Take the extreme emotivist view which says that in uttering “That was wrong” I am not asserting anything but only expressing my own distaste or my disapprobation of the act in question. That view can really give no account of the validity of the inference at all. In saying “If that act was wrong then—” my “If” does not signify that I am somehow hypothetically or conditionally expressing a feeling. There is no such thing as hypothetically expressing a feeling. Of course, I can say “If I feel such-and-such

then ——” and then reflect on what follows from my having a certain feeling, or what would be true if I had one. I can also draw conclusions from the supposition not that I have a certain feeling but that a certain feeling has been expressed. But the antecedent in all those reasonings would be a straightforward factual proposition about feelings or the expression of feelings. They would not be mere expressions of feeling which are neither true nor false.

Other kinds of nonpropositional theory are more complicated, but they all face similar difficulties. Some hold that to make a moral judgment is not to say anything true or false but to prescribe a certain course of action in the way that imperatives order or demand certain courses of action. But still there is a difficulty about how one can hypothetically prescribe or recommend something. One can certainly prescribe or order something that is hypothetical or conditional —“If you go out, shut the door after you.” But that is an order to do something if certain conditions are fulfilled. The imperative does not appear as the antecedent of a conditional proposition. Moral judgments like “That was wrong,” it appear, do occur as the antecedents of conditional propositions, and inferences are validly drawn from them. But if they are prescriptions, it would seem that in such positions they serve to issue prescriptions only conditionally. And what could that be? It would not be entertaining the hypothesis that a certain prescription has been made or that a certain course of action has been recommended. Those are both straightforward factual propositions which are either true or false. They can easily be embedded in other sentences. But how could a prescription itself be embedded in a conditional sentence? It seems that it would have to be something like a hypothetical issuing of a prescription. But there is no such thing.

Another type of view, perhaps closest to what Hume says about causal necessity, is that in making moral judgments we do take ourselves to be expressing beliefs which are either true or false,

but we are deeply confused and mistaken. What we are really doing is projecting something we feel when perceiving or thinking about an action onto that action itself and mistakenly supposing that it “objectively” resides there. “The mind,” Hume says, “has a great propensity to spread itself on external objects, and to conjoin with them any internal impressions, which they occasion, and which always make their appearance at the same time that these objects discover themselves to the senses.”⁸ In making moral judgments we think we are ascribing moral characteristics to the acts we observe; we treat our moral views as if they were, so to speak, propositional, but in fact they are mere projections. We “gild” or “stain” the facts with our feelings, but all that is strictly true in what we say is the purely factual, nonevaluative content to which something in the value-free world could correspond.

This kind of view seems to me to serve the interests of “subjectivism” best. But so far it gives no account of what our making a moral judgment really amounts to. It does not explain what we are saying when we say or believe that a particular act was wrong. We are said to take something we feel and project it onto the world, believing it to be a property belonging to things that exist there. But how do we do that? We do not think that objects and events in the world actually have the very feelings that give rise to our own “pronouncements.” The most that could be said is that we ascribe to things in the world, not the feeling itself, but what the feeling is a feeling of —that very feature that we are aware of “in our own breast.” In the case of causality Hume thinks we get what he calls an impression of necessity, and it is that very feature —necessity —that we ascribe to the connections between some of the events we observe. In the case of color it is, say, greenness that we perceive and then project. What is the corresponding feeling or impression in the moral case? What we think or judge is that the act was bad, or vicious, or wrong, so it would

⁸ *Treatise*, p. 167.

seem that it must be a feeling of badness, or vice, or wrongness. But what *is* such a feeling? The question for “subjectivism” is whether and how we can understand those particular feelings that it says either generate or are referred to in our moral judgments or opinions.

Hume calls the feelings in question sentiments of approbation or disapprobation. But what makes a feeling a feeling of disapprobation or disapproval? Not just any bad or negative feeling will count. To disapprove of something is to think it bad, to make an unfavorable evaluation of it. So a particular feeling will be a feeling of disapproval only if it is generated by or suffused with the thought that the thing in question is bad. But that is precisely the evaluative thought that the theory is trying to account for. It must explain how we can think something is bad or wrong without itself attributing badness or wrongness or any other evaluative feature to anything.

This same difficulty faces other versions of “subjectivism” in which feelings are said to play an essential role in moral judgments. If my moral judgment is a report that I have a certain feeling, or that all human beings would get a certain kind of feeling under certain conditions, the kind of feeling in question must be identified before we can know what is being said. Not just any feeling will do. Hume says that the feeling arising from virtue is “agreeable,” and the feeling of vice is “uneasy” or unpleasant, but in saying that an act is wrong, even if I am indeed saying something about how people do or would feel, I am not saying only that they would get unpleasant or disagreeable feelings from the act. They might get unpleasant feelings from something they eat, but that would not make what they eat bad, or vicious, or wrong. So we still need some explanation of what it is to think that something is bad, or vicious, and some account of how we can intelligibly attribute such thoughts or attitudes to people.

I do not mean to suggest that, as things actually are, there is any special difficulty about our doing that. We often agree in our

moral assessments of particular acts and in many of our more general evaluative opinions. We come to share values, when we do, by growing up and living in a culture in which they are endorsed and acted on. We recognize the badness of certain acts, and we recognize that other people have beliefs or reactions that are appropriate to the badness of the acts we all observe. Their responses count as disapproval because they involve the thought that the acts are bad —a thought which we know to apply truly to just such acts as these. Our ascriptions of evaluative attitudes or feelings to human beings go hand in hand with our ascriptions of value to things and actions in the world.

I need not share all those moral assessments that I can correctly attribute to others. I can recognize that others think that a certain sort of thing is bad even if I do not think the thing is bad, because I too can have that same thought about other things. I do not have to agree in each particular case, any more than I must agree with someone else's judgment about the color of something in order to attribute a belief about or a perception of color to that person. Knowing that a blue light is shining on a white wall, I will know that a person looking at it sees blue and, if he doesn't know about the light, that he also believes that the wall is blue. I know the belief is false, but I can attribute that belief to him. I can do that because of my own general competence in the language of color and my knowledge of what colors things are in the environment. Similarly, if I do not agree with a person's evaluative judgment, I can still correctly attribute it to him and understand what it is for him to hold that view, because of my own general competence in the language of evaluation and my knowledge of the evaluative features of the environment —what things are good or bad, better or worse than others.

The traditional theory of simple and complex ideas was a way of accounting for the possibility of false belief or of a lack of correspondence between people's ideas and the world. But that theory seems no more plausible here as a way of understanding the

possibility of evaluative thought in general than it seemed in the case of color. Perhaps some particular evaluative concepts can be defined in terms of others, but surely we cannot expect all evaluative notions to be reduced to terms that are not evaluative at all. There are no simple nonevaluative “elements” which could somehow be combined in thought or experience to give us the idea of value, and hence the possibility of evaluation, in the first place. This irreducibility is one of the few things on which most modern moral philosophers would seem to agree.

There is no question that we do make moral judgments or evaluate things or states of affairs, and that we do attribute such judgments, or reactions involving such judgments, to others. The question is not whether we all do it in real life. The question is whether someone who consistently holds to the “subjectivity” of all values could do it. Could someone make sense of the idea of there being feelings or attitudes of disapproval, say, if that person did not also hold the view that certain kinds of acts are bad, or wrong, or worthy of disapproval? What made it seem possible in the case of color was the thought that perceptions could somehow be directly recognized as intrinsically of a certain specific kind —that we can simply read off from our perceptions themselves what features they are perceptions of —whatever we take the world to be like, whether we think it contains any colored things or not. I think there is a tendency to rely on a similar thought in the case of values. We are thought to be able to recognize what we feel simply by feeling it, by being aware of some felt feature in our experience, whatever we take the world to be like, whether we attribute any negative or positive value to anything in the world or not.

I have already suggested why I think that sort of view could not be right even about perceptions of color. I do not think its prospects are any better in the case of values. Even supposing that we could isolate in our experience some feeling or attitude or response which plays an essential role in moral or evaluative judg-

ment, there would still be the problem of what that feeling could be said to be like. It would have to be identified and classified only in terms that are somehow immediately available to consciousness, not in terms of any evaluative judgments that define or accompany it. It would have to be the kind of feeling or response that a person could have without having any moral or evaluative opinions at all.

This would have the consequence that the only materials available to us for understanding what appears to be evaluative thought and for seeing how it figures in human action and human social arrangements would be simple, isolated feelings with no evaluative content. They might be such things as pleasant or unpleasant sensations, feelings of satisfaction or dissatisfaction, or simple likes or dislikes. Or, moving away from feelings, they might be such things as basic unmotivated desires or wants or preferences, or even more indiscriminately, those all-purpose motivators called “pro-attitude” and “con-attitude.”

Even such apparently scaled-down materials are not necessarily on the nonevaluative ground floor. One can feel pleased that justice has been done, for example, and if that is a feeling of pleasure it is still not independent of its evaluative content. Someone who did not think that justice had been done could not have such a feeling. And liking something or somebody can be a matter of thinking well of the thing or person, and that again has an essential evaluative component. Even wanting, or preferring, or simply being for or against a certain thing can also be an evaluative attitude or state. Preferring that virtue be rewarded, for example, or being for a just solution, or being against the unjust acts of one’s government —these are moral attitudes and not simple motivating feelings or wants that might rise up in a nonevaluating agent. It is not clear to what extent there could even be such a thing as a nonevaluating agent —at least, a human agent.

I have suggested that it is the thought that values and evaluation would not otherwise be intelligible that can lead to the idea

that they must ultimately be explained only in terms of likes and dislikes, pleasures and pains, or basic human desires. Perhaps something like that is what finds expression in the popular half-thought that morality is after all just a matter of what people want, or what they like or don't like. Or worse still, the thought that it is just a matter of whose likes and dislikes are going to prevail. And now there is the view, in the United States at least, that morality is itself just one among a great many "special interests" that have to be accommodated in society. People are thought to be just pushing their own personal interests or seeking their own "gratifications" in one way or another, and the "morality lobby" is encouraged to fight it out with the military, the corporations, the doctors, the judges, and so on.

Whether such views are really derived from the position I am calling "subjectivism" I don't know. But it does seem to me that to hope that the feelings or attitudes essential to evaluative judgment can be identified and understood neutrally, on the basis of some intrinsically felt quality alone, would be disastrous for making sense of what is, after all, a fundamental aspect of human life. If it is only our very engagement in a set of values that makes it possible for us even to recognize the phenomenon of evaluation, the demand of evaluative neutrality would have the effect of denying or obliterating the very phenomenon we want to understand. If engagement or participation is essential, we can never get ourselves into a position to discover that all values are "subjective," that the goodness or badness of something is not part of the way things are.

II

I have been trying to identify some of the difficulties in "subjectivism" and to draw attention to what looks like a serious obstacle to our ever even arriving at that conception of the world. If we tried to adhere strictly to what "subjectivism" requires of us, it seems that we could not consistently or coherently come to think

that it is true of our evaluative thought. But even if that is so it might not seem like much of a threat. Why can't we simply abandon "subjectivism" and look elsewhere for a more satisfactory understanding of evaluative thought and its role in our lives? I now want to suggest why I think that will not be easy.

"Subjectivism" in one form or another appears to be the inevitable result of our trying to understand ourselves in a certain way, and the goal of understanding ourselves that way is not easily abandoned. It seems impossible to conceive of a better way of understanding how human beings work. The model goes back to one of the greatest achievements of the Enlightenment —the idea of a "science of human nature." It was to be a thorough, systematic investigation of the principles of human nature that would eventually explain every aspect of the personal, social, cultural, and political life of human beings. That the proper aims of individual human beings and the best social and political arrangements among them should somehow be determined in the light of truths about their nature and the world they live in was not in itself a new idea. What was unique to the Enlightenment was an open-ended curiosity about what that human nature is really like and, most important, an empirical, secular idea of how it is to be known.

If human beings were endowed by a supernatural power with certain capacities and goals and were placed in a specially ordained position in the world, that would obviously limit both the content and the justification they could find for the beliefs they arrived at and the goals they aspired to. The source of any knowledge they acquired or any values they pursued would then lie in something other than those human beings themselves and the familiar world they could know they inhabit. Their behavior in pursuit of that knowledge and those values would therefore not be fully understandable to them by human intellectual means alone. There could be no properly scientific understanding of human behavior or human life.

The impressive growth in understanding the inanimate physical world, culminating in the work of Isaac Newton, was the decisive step toward the downfall of that traditional picture. The mathematical science of nature was seen to rest on no supernatural hypotheses and to proceed carefully no further than experience, and solid reasoning based on that experience, could take it. It does not matter now whether that really was a correct perception of the revolution in physical science or not. In any case it served as the source of the Enlightenment ideal of understanding what we might call the human world through the application of just such a broadly scientific enterprise to the study of human nature and human life. It put human beings at the center, and it insisted that the way things should be in any human world must be based only on what human beings can reliably find out about themselves and the natural observable world they live in.

One of the earliest and still one of the best statements of this goal is Hume's *A Treatise of Human Nature*. It was written in the belief that "all the sciences have a relation, greater or less, to human nature. . . . Even *Mathematics, Natural Philosophy, and Natural Religion*, are in some measure dependent on the science of MAN; since they lie under the cognizance of men, and are judg'd of by their powers and faculties." The aim of the book was to lay the foundations of a genuine science of man that was to be like nothing that had gone before.

Hume had found that the "moral Philosophy transmitted to us by Antiquity," like ancient "Natural Philosophy," was "entirely Hypothetical, & depend[ed] more upon Invention than Experience. Every one consulted his Fancy in erecting Schemes of Virtue & of Happiness, without regarding human Nature, upon which every moral Conclusion must depend." ¹⁰ His plan was to appeal to nothing but experience—to "introduce the experimental method

⁹ *Ibid.*, p. xix.

¹⁰ *The Letters of David Hume*, vol. 1, ed. J. Y. T. Grieg (Oxford: Clarendon Press, 1969), p. 16.

of reasoning into moral subjects.” He would take human beings as they are, he would observe them “in the common course of the world, . . . in company, in affairs, and in their pleasures,”¹¹ and he would try eventually to explain all the overt and hidden richness of their behavior, thought, and feeling. That would require the discovery of general principles of human nature, but the only possible access to such principles and explanations would be judicious generalization from whatever we can find out by “the cautious observation of human life.”¹² Where our best experience remains still silent on some question of human nature or human destiny we must willingly confess our ignorance and perhaps try harder to discover how we work and what the world holds in store for us. But we must not let the natural human desire for some answer or other lead us to invent comforting stories for which we can honestly find no support in experience, or impose such conjectures or hypotheses on the world and then base our beliefs and behavior on such creatures of the imagination.

The positive project of a naturalistic study of human beings is familiar to us today in what we call the social sciences. They are such a pervasive and powerful feature of the modern world that it is sometimes difficult to remember that they have not always been with us. Ideally, they promise a dispassionate, scientific understanding of thought, feeling, and behavior in every area of human life. The information they would provide is to be the basis of all personal, social, and political organization and improvement, just as the physical sciences provide the facts and theories that engineering and technology then make such spectacular use of in the purely inanimate domain. By now we are used to a division of the study of human behavior into different, highly specialized fields such as economics, psychology, and sociology. Hume probably never envisaged the technical professionalism of today. But

¹¹ *Treatise*, p. xxiii.

¹² *Ibid.*, p. xxiii.

the idea that human beings can be studied and understood in this way is an Enlightenment idea. It is just what he had in mind.

Hume's enthusiasm for the idea went beyond the social benefits to what he saw as the directly cognitive or intellectual payoff of a science of man. Since all sciences—even mathematics and physics—fall under the cognizance of men, he thought it was “impossible to tell what changes and improvements we might make in these sciences were we thoroughly acquainted with the extent and force of human understanding, and cou'd explain the nature of the ideas we employ, and of the operations we perform in our reasonings.”¹³ These hopes have scarcely been vindicated. It is difficult to point to concrete advances in the sciences that have been generated by Hume's own epistemological theory of our ideas and reasonings, and presumably few today would think of looking to psychology, say, or sociology, as a possible source of breakthroughs in mathematics or physics. But even without the hope of such direct scientific consequences, the Enlightenment idea of the study of human nature still serves to determine our culture's conception of understanding ourselves. It remains difficult to imagine any other way of getting the kind of understanding of ourselves that we seek.

The goal is an understanding of human nature. The method is to study human beings in interaction with the world and thereby to explain how they come to think, feel, and act as they do. But not just any story—even any true story—about the relation between human beings and the world will give us what we want. For instance, we are interested in how people come to believe and know what they do about the world around them, but we would not be satisfied with the obvious truth that they learn to think and speak and they come to know things about the world by seeing and touching things and in other ways perceiving what is true of the things around them. That is all true, but it does not

¹³ Ibid., p. xix.

explain what we want to understand. Similarly, we are interested in how people come to value things and why they endorse the particular values that they do, but we would not be satisfied with the obvious truth that they grow up and are socialized into a particular culture and come to accept many of the value beliefs current in that culture. That again is certainly true. But such general truths about human beings and human life do not explain what we want the philosophical study of human nature to explain.

It is not just a question of detail. Even a full, detailed story of how a particular person or a particular group comes to know or to value a certain sort of thing would not satisfy us. We want to understand certain pervasive or fundamental aspects of human life *in general*.

Morality, for example, is a quite general phenomenon which seems distinctive of the human species. No other creatures seem moved by considerations of good and bad, right and wrong. To understand that aspect of human life, then, would be to understand how there comes to be such a thing as morality at all, how it works, and what makes it possible. Our possession of an elaborate conception of the world we live in and its history, how it works, and how it affects us, is also something a science of human nature should be able to explain. We want to know how there comes to be any such thing at all. The same is true of our beliefs about the colors of things in particular, or all our beliefs to the effect that certain things are causally connected with other things. The question in each case is how we come to have any thoughts or beliefs or responses of that general kind at all.

There is implicit in this kind of question a certain idea of how best to answer it. It seems to be the only way in which it could be answered. Obviously, if human beings come to act or think or feel a certain way only after interaction with the world around them, there must be something about human beings —something about what we are like —and something about what the world is like

which combine to produce in human beings the way of acting or thinking or feeling in question. Even if facts of the world affected us directly by simply impressing themselves on our minds whenever we opened our eyes or ears, there would still be something true of us that was partly responsible for our getting all the beliefs we get. It would be because we are capable of passively receiving information about the world in that way. So it seems that any explanation of distinctive, pervasive features of human life would have to be a two-part explanation. It would involve an “objective” factor — what the world is like, or how things are independent of us — and a “subjective” factor — what we are like, or how things are with us.

If both factors are always present, if human beings themselves always play some role in acquiring their conception of how things are, then the study of human nature will naturally take the form of asking how much, and what, the human subject does contribute. That would be to isolate and identify those elements of human nature that are responsible for our conceiving of and responding to the world in the ways we do. How much of what we think or feel about the world is due to us, to the way we are, to the “subjective” factor, and how much is due to the way things are independently of us, to the “objective” factor? The intellectual goal expressed in terms of this bipartite conception serves as our model for understanding ourselves, for seeing how we work, for identifying what is distinctively human.

Hume’s Enlightenment project of a “science of human nature” embodied just this conception of how to understand ourselves. It was the search for, among other things, “principles of human nature” — those features of human minds and sensibilities that are responsible for our thinking, feeling, and acting as we do. Hume sought those “principles” which he called “permanent, irresistible, and universal.” They are the “foundation” of all our thoughts, feelings, and actions in the sense that, as he put it, “upon

their removal human nature must immediately perish and go to ruin.”¹⁴ They will be those “principles” that are at work in anything recognizable as human life as we know it —ways of thinking and acting without which nothing would think or act as humans do. They are to be identified by asking what must be true of human beings as we know them in order for them to think and feel and act in the ways we know they do.

This Enlightenment project of isolating and identifying the “subjective” factor in human thought and experience is by no means only a thing of the past. We find it in any philosopher who would distinguish in general between the “given” and the “interpretation,” between the “data” we receive and the construction we put upon them, or between the “flux of experience” and the “conceptual scheme” we impose upon it to make sense of our experience and to learn from it. In our own day, for example, W. V. Quine in his *Word and Object* has put his task almost mathematically as follows: “we can investigate the world, and man as a part of it, and thus find out what cues he could have of what goes on around him. Subtracting his cues from his world view, we get man’s net contribution as the difference. This difference marks the extent of man’s conceptual sovereignty.”¹⁵

It is not just a matter of how the world affects us. We know that something outside us acts on us; our cognitive and affective responses are caused by something in the world. So in that sense we know that there is something or other in the world which, in conjunction with facts about us, makes us think and feel as we do. But it is not merely a question of causation. We are interested in a more complex relation between our beliefs and responses and the world. We want to know the nature of whatever causes there are. Is there anything in the world that not only causes but somehow matches up with, or corresponds to, or is adequately repre-

¹⁴ Ibid., p. 225.

¹⁵ W. V. Quine, *Word and Object* (Cambridge, Mass.: Technology Press of MIT and John Wiley & Sons, 1960), p. 5.

sented by, those things that we think and feel? Is the world, or does the world have to be, anything like the way we think it is, in order for it to have given us the thoughts and feelings about it that we have? What is there, if anything, that renders those thoughts or responses true? Or are they merely “subjective” responses with nothing corresponding to them in the “objective” world? And with these questions we have arrived at the inquiry that I have already suggested seems so easily to lead to “subjectivism” with respect to many of the things we believe.

It leads easily to skepticism about the world too. Descartes’s evil demon represented the threat that the way things are independently of me might be extremely different from the way I take them to be. If the demon exists, he alone exists beyond me, and his clever machinations make me think that I live in a world of earth and water, trees and buildings, and other people with human bodies like the one I think I’ve got. He gives me such thoughts and beliefs, so they are produced by something “objective” and independent of me, but there would be almost nothing in that world corresponding to any of those thoughts. They would almost all be false. The challenge of skepticism is to show how I know that I do not live in such a world. Once we see human knowledge as a combination of an “objective” and a “subjective” factor in this way, and we acknowledge the possibility of a largely or even entirely “subjective” source for most of our beliefs, it seems impossible to explain how those beliefs could ever amount to knowledge or reasonable belief.

But our concern here is not the epistemological question of knowledge or the reasonableness of our beliefs. The “subjectivism” I have been describing results from a metaphysical project which relies on some unquestioned knowledge of the world. It subtracts from some of our beliefs, not the causes that produce them, and not necessarily our warrant for accepting them, but rather their correspondence with anything that holds in the way things are independently of us. Whatever support we might have

for them therefore cannot be “objective.” It cannot be based in anything that is to be found in the way things are. The same bipartite conception of the ultimate source of all our beliefs and responses is at work in this project. Everything is to be assigned either to an “objective” or to a “subjective” source.

It is because we want to understand human values or human valuing in general in terms of this bipartite picture that I think we are inevitably driven toward “subjectivism.” To understand the general phenomenon of human evaluation we want to understand the content of such evaluative thoughts — what exactly is being said or thought when someone says, “That was a vicious act” or “It is better to comfort someone than to kill him.” To ask what is involved in such a thought is to ask what, if anything, could make it true, or false. If we answer, simply, that it is the goodness or badness of things that makes such remarks true or false, we feel we are not really explaining the content of those evaluative “pronouncements.” In saying that it is true that an act is vicious if and only if the act is indeed vicious we are making use of the idea of viciousness, but we are not explaining it. We feel we will understand evaluation only when we know, as it is often put in philosophy, *what it is* to be vicious, or bad, or wrong.

When Hume investigated that fundamental feature of human life which he called our “reasoning from causes to effects,” our getting beliefs about the necessarily causal connections between things, he did not restrict himself simply to identifying the circumstances in which we get such beliefs. If he had, he could have said that a belief that two things are causally connected arises in us whenever we are presented with instances in which one thing causes another. Even if that were true, it would not help us understand human thought about causality in general because it does not explain what it is for one thing to cause another. It was because Hume wanted to explain not just the origin but also the contents of our causal beliefs that he was able to find nothing in the world corresponding to their special claim of necessity. Neces-

sity, then, could be only something “subjective,” or nothing at all. It was the very desire for a completely general account that would explain how any thought about necessity is possible at all that led to “subjectivism” about necessity.

If the contents of our evaluative beliefs are going to be explained in terms of what is or could be the case in the world, it does seem that we will have explained evaluative thought in general in that way only if whatever we find in the “objective” world to explain it is something nonevaluative. Otherwise, there will still be some evaluative content that will not have been explained. And since the nonevaluative always falls short of exhausting the special, apparently evaluative, content of our thought, the source of that special evaluative element will inevitably have to be located in us, not in things as they are independently of us. It will be assigned to the “subjective” and not the “objective” factor. All our value beliefs, then, insofar as they are really evaluative, will be “subjective” and will not assert anything that is or even could be true of the way things are.

In the first lecture I tried to cast doubt on our ability to carry out this metaphysical project of conceiving of a world in which we have genuinely evaluative beliefs or attitudes while holding that none of them is true of the world. It requires detachment from, or nonendorsement of, the contents of all our evaluations. And it requires an acknowledgment of the fact that we nevertheless do make such evaluations. But the fulfillment of either one of those two requirements threatens the possibility of fulfilling the other. Without ascribing value to things in the world, and hence holding evaluative beliefs of our own that we take to be true, it is difficult to see how we could interpret and hence understand other people, as well as ourselves, as holding any evaluative beliefs at all. We would not have what it takes to see the world as containing genuine evaluations — any thoughts to the effect that something is good or bad, or better or worse than something else. If we do find thoughts in the world which we understand

to be genuinely evaluative, then it seems that we must already hold certain evaluative beliefs or opinions of our own. But then our disengagement from all values would have been abandoned. We would be taking certain value sentences as true.

I did not try to prove any of that. I said what I could to make it plausible. Without offering further argument for it here, I want to examine some of its consequences. Suppose, as I have been suggesting, that there is simply no understanding of evaluation in completely nonevaluative terms. It has long been accepted that there is no hope of strictly *defining* evaluative notions somehow in purely nonevaluative terms. That is, perhaps, the real lesson of G. E. Moore's misnamed "naturalistic fallacy," which preoccupied moral philosophers for so long.¹⁶ I am suggesting, not an obstacle to definition of the contents of our evaluations, but rather an obstacle to our even understanding or acknowledging the phenomenon of human evaluation at all. Evaluation is a fact of human life—something that human beings do—which it seems we cannot acknowledge without our also engaging in the practice ourselves. And that means that we cannot even understand that it is going on without our being prepared to take certain evaluative beliefs or "pronouncements" to be true.

Just suppose for a moment that that is right. Now if it is also right (as I think it is) that accepting "subjectivism" about values would have to involve our disengaging from all values while still making sense of evaluative thought, then we would not really be able to accept a "subjectivist" picture of human values. We could not coherently get ourselves into the position of discovering that none of the evaluative beliefs or attitudes of human beings corresponds to anything that is so in the way things are. But then, if "subjectivism" is the inevitable outcome of trying to understand human values in terms of that traditional bipartite conception of human beings and their relation to the world, it would follow that

¹⁶ G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1956), chap. 1.

we cannot really carry out the Enlightenment project of determining the “objectivity” or “subjectivity” of values in general. The detachment or disengagement we would need would rule out the very understanding that we seek. That would be disturbing, and dissatisfying, given the natural appeal of that picture. It seems like a perfectly comprehensible intellectual goal—in fact, the very model of what it would be to understand general aspects of human nature in the right way. But it would be unattainable. We could never fully understand ourselves in that way.

This is not to say that we could not study the phenomenon of human evaluation, and indeed human values themselves, and learn much more than we now do. There is a great deal that we do not know and should be trying to find out, not only about what things are good and bad, and why, but also about how people acquire the values they do. How does it happen that an infant who comes into the world with needs and impulses and a native set of behavior patterns comes by the time it is an adult to possess a complex set of evaluative beliefs and responses? How do adults with firmly held evaluative opinions about certain matters come over the course of time to change them? And how can we arrange things in society so that people on the whole make such changes in the direction of more informed, more considerate—in a word, better—evaluative attitudes? There is no answering such questions in the abstract. It obviously depends on the particular people and the particular culture and on countless other factors in ways we still do not understand very well. But such questions, however complex, can be answered—or at least progress can be made.

Any study of human socialization or human development along these lines would be a study of how a human being or a group of human beings gets absorbed into a culture whose members already have some values or other, or how the possession of one set of values gets transformed into possession of another. It would explain at most the transmission of values, perhaps even the transmission of the very idea of value, from those who have it to those

who do not. But to explain how something is transmitted or changed is not necessarily to explain what it is that is transmitted or changed. It is not necessarily to explain what those attitudes are, or what it is to hold them. What is it to think that deliberate killing of a human being is a very bad thing, or to think anything evaluative about anything? This is a philosophical question about human valuing as such.

If that is what we want to understand about ourselves or about human nature in general, then the metaphysical project I have been describing seems inevitably to come into play. This is what leads us to “subjectivism.” We want to understand the nature of any evaluative thoughts or attitudes. We ask what their special content is, what is really being thought. And that first takes the form of asking what would be so if they were true. In trying to answer that question, either we merely repeat the thought — “Killing a human being is bad” is true if and only if killing a human being is bad — and so we do not feel we are explaining it, or we try to express its content in other terms that reveal in some illuminating way what is really being said. If those further terms are still evaluative, we will not feel that we have explained what it is for any evaluative thought or attitude to be true; we will simply have exchanged one such thought for another.

So if we are going to make any progress in explaining the evaluative as such, we will either say that having what we call an evaluative attitude or opinion is not really a matter of thinking something to be true — but instead is expressing a feeling or issuing a prescription or making a recommendation or some such thing — or we will say that it is a matter of our thinking true something that is really nonevaluative and so could hold in the “objective” world — perhaps something about nonevaluative feelings that we and others do or would feel under certain conditions. Each of these alternatives is a version of what I am calling “subjectivism”: there are no “objective” evaluative facts or states of affairs. And each of these alternatives appears to deny or obliterate

the very phenomenon we set out to understand. That looks like the inevitable outcome of our trying to understand human evaluation in general.

It is this very desire to explain human evaluation in general that seems to preclude us from invoking unexplained evaluative truths or states of affairs in any account of the special contents of human evaluations. Thus do we inevitably banish genuinely evaluative facts from any world in which we can make sense of what seems to be evaluation. But in trying consistently to adhere to nothing more than that shrunken conception of what is really so we would fail to make sense of the idea that human beings have such things as evaluative thoughts, attitudes, or responses. We would lose those very attitudes that “subjectivism” about values claims have nothing corresponding to them in reality but are nothing more than our “subjective” responses to an “objectively” value-free world.

I would draw here on the parallel I see with the case of colors. If we did not make categorical ascriptions of colors to things around us we could not acknowledge the existence of such things as perceptions of colors or beliefs about the colors of things on the part of human beings. We could not conceive of the world as containing those very perceptions and beliefs that “subjectivism” about colors claims have nothing corresponding to them in reality and are nothing more than our “subjective” responses to an “objectively” colorless world.

I want to say more about what this idea amounts to and exactly what it implies about our understanding of values and colors, and what it does not. It says in its strongest form that we cannot think of a world in which people perceive particular colors or believe that things are colored without ourselves being prepared to ascribe color categorically to things in the world. We cannot understand human beings to have evaluative opinions or attitudes to the effect that such-and-such is good or bad without ourselves sometimes recognizing the goodness or badness of certain things. And in

making those ascriptions of color, or of value, we are taking certain things to be true. We take it to be part of the way things are, for example, that grass is green, or that the deliberate killing of a human being is a very bad thing. Our engagement with, or endorsement of, aspects of the world of those general types is required for our ascribing to human beings beliefs or attitudes with those types of contents. What we take to be facts of the world are implicated in our making sense of thoughts of the world. The two cannot be pried apart completely.

I am here endorsing particular instances, having to do with colors and with evaluation, of what appears to be a quite general fact about our understanding one another in the ways we do. There are conditions of the successful ascription to human beings of beliefs, attitudes, perceptions, or feelings with specific contents. This is something that I believe lies at the heart of Wittgenstein's later work. Donald Davidson has stressed its importance in what he calls "radical interpretation."¹⁷

If we are to interpret someone as believing or perceiving or feeling some particular thing or as having a certain specific attitude, we must somehow connect those specific psychological states we are attributing to that person with facts or events or states of affairs in the world that we take them to be about. If we ourselves had no opinions about what is so and what other people are most likely to be attending to in the environment, we would be in no position to attribute any beliefs or perceptions or attitudes to them at all. We interpreters and ascribers of beliefs and other psychological states must therefore be engaged in the world and take certain things to be true of it if we are ever going to attribute psychological states to anyone. And we have no choice but to ascribe to others, at least in general, beliefs in and perceptions of and attitudes toward some of the very things we ourselves take to be true of the world. We cannot make sense of other people

¹⁷ See, e.g., Donald Davidson, *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984), essays 9–12.

as believing something we know to be obviously false unless we have some explanation in the particular case of how they come to get it wrong. And that explanation will work only if we understand them to share in common with us other beliefs and attitudes in the midst of which their particular, localized error (as we see it) can be made intelligible.

This still leaves considerable room for difference or disagreement. Whole areas of belief or perception might be found to diverge if there remains enough overlap to serve as the shared base of what the interpreter could then see as the others' deviance. But it seems to me, as I have been suggesting, that the deviance could not go as far as the interpreter's finding that others had perceptions of color and beliefs about the colors of things which he held did not agree with anything at all that he took to be true in the world. Nor could he find that they had evaluative beliefs or attitudes about the goodness or badness of things none of which he shared because he had no evaluative beliefs or attitudes at all.

Davidson has sometimes drawn from his main claim about interpretation the conclusion that most of our beliefs must therefore be true.¹⁸ They must be, if we can even understand the fact that we and other people have any beliefs at all. And this seems to imply that the truth of the majority of our beliefs is a necessary condition of our having them —that if we have any beliefs or attitudes at all, the list of sentences which state the contents of those beliefs or attitudes will contain mostly truths. That would connect our beliefs necessarily with the way the world is. If I am right to apply this thesis about interpretation to color beliefs in particular, that would imply that most of our color beliefs are true. It would mean that we are getting the colors of things, on the whole, right. We can't help it. And applied to our evaluations it would mean that, in general, our beliefs about what is good and bad, better and worse, are true. On the whole, the things that we

¹⁸ Donald Davidson, "A Coherence Theory of Truth and Knowledge," in *Kant oder Hegel*, ed. D. Henrich (Stuttgart: Klett-Cotta, 1983).

think are good, or bad, really are good, or bad. Of course, in the case of values there appears to be much less widespread agreement, so it is not easy to speak without further qualification of what “we” believe, what “our” evaluative beliefs are. But despite that apparent lack of agreement, there must still be enough common ground somewhere among all those who hold evaluative beliefs to make possible the ascription of such beliefs to them. That common core, or at least the major part of it, is then said, on Davidson’s view, to be true.

This is not a conclusion I wish to draw — at least not if it is taken as a defense of the “objectivity” of colors or values as opposed to their alleged “subjectivity.” If we knew, by this kind of transcendental argument, that most of our beliefs had to be true, and in particular that most of our color beliefs are true, and that most of our evaluations are true, we could easily be led to ask, in the spirit of the traditional metaphysical project, how that could possibly be so. How could it be that our believing what we do requires that the world should be a certain way? This will make us look once again at the contents of those beliefs. What exactly are we believing when we believe that, say, grass is green, or that the deliberate killing of a human being is a bad thing? And what is it to believe such things? We want to understand such beliefs in general. Just as that quest drove us toward “subjectivism” earlier, so it would drive us toward “subjectivism” again. There seems to be no other way to account for the necessary connection that would have been proven to hold between the body of our beliefs as a whole and their truth.

This is just the position of Kant’s *Critique of Pure Reason*, perhaps the greatest attempt there has ever been to prove that the truth about the way the world is cannot come apart in general from our thinking and perceiving in the ways we do. Kant thought there were necessary conditions of the possibility of all thought and experience, and that not all those necessary conditions are themselves just further thoughts or beliefs. They include as well

many nonpsychological truths, so not only must we think a certain way, but the world independent of us must be a certain way, in general, if we are even able to think of or perceive anything at all.

Kant saw that some philosophical theory was needed to explain this necessary link between thought and experience and the world, and his explanation was the theory of transcendental idealism. It was the only explanation he thought there could be. We can know that the world in general must conform to our thinking and perceiving in certain ways because our being able to think and perceive what we do actually “constitutes” the world that we perceive and believe in. The price of showing that our thoughts and perceptions and the truth of their contents cannot come apart in general was that the truth of what we believe about the world somehow consists in our having the kinds of thoughts and perceptions that we do. The world turns out to be dependent on our thoughts and perceptions in some way after all. That is a form of idealism, which is one variety of what I am calling “subjectivism.”

If we ask in a similar vein how our color beliefs, or our evaluative beliefs, could not fail to be true, a more particular version of that same idealism or “subjectivism” will seem like the only possible answer. What it is for color judgments to be true, for there to be a world of colored objects, it would say, is just for human beings to agree for the most part in their ascription of colors to things. There might be considerable disagreement in particular cases, but on the whole there would be nothing more to things’ being colored than human perceivers agreeing in general in the perceptions they have and the judgments they make about the colors of things. Similarly, for things to have value, and to have the particular values they have, would simply be for human beings to agree in general in their ascriptions of value to things. Again, there is room for wide disagreement and uncertainty, but on the whole the truth of value judgments would amount to nothing more than agreement, or the possibility of agreement, in human beings’ evaluative beliefs.

This is clearly just “subjectivism” approached from a different direction. In terms of the traditional dichotomy, it locates the source of the truth of color judgments and of value judgments on the “subjective” side. Or rather, like all forms of idealism, it in effect collapses what was originally thought of as the “objective” into the “subjective.” The facts of the “objective” world that make our color judgments or our value judgments true would be facts only about us, about what we say and do, and not about an independent world that we say those things about. The only form of so-called “objectivity” granted to those judgments would be intersubjective agreement. And understood in terms of the traditional dichotomy that is not really “objectivity” at all.

This is precisely the kind of view of the world that I have been suggesting we can never reach. Describing it this way perhaps brings out why. We cannot make sense of the idea that the truth of judgments of a certain kind amounts to nothing more than human agreement with respect to the contents of those very judgments. If human beings agree in certain judgments, there must be something they agree about. We must be able to make independent sense of their making such judgments in the first place if we are to find anything for them to be in agreement about. That is the point of Davidson’s requirements on interpretation. The content of the judgment must be identifiable independently of the fact that the judgment is made. And that is why our taking certain things to be true of the world must be involved in interpretation from the beginning.

Davidson himself would make no appeal to idealism or any form of “subjectivism” to explain why he thinks most of our beliefs must be true. It comes from the conditions of interpretation alone. But making sense of what people are saying and doing, and ascribing various psychological states to them, is something that we human beings do. The conditions of interpretation or understanding are conditions of our doing something, or our succeeding in doing it, not simply conditions of something’s being so

independently of our efforts to understand. If our taking certain things to be true is a condition of ascribing to people beliefs and perceptions and attitudes with specific contents, then there will necessarily be considerable agreement among us. But that does not strictly imply that what we largely agree about must be true. The truth of something does not in general follow from the fact that some or many or even all human beings agree about it. Nothing about the conditions of interpretation can obliterate that fact.

Of course, if we all agree about many things then we will *regard* them all as true. We will hold, of those things that we all agree about, that they are all true. But it still does not follow from our acknowledged agreement that they are all true, even if we insist that agreement is indeed necessary for interpretation and mutual understanding. So even if we must all share certain evaluative (or color) judgments if we are to see ourselves as having any evaluative (or color) opinions at all, I do not see that the truth of any evaluative (or color) statements themselves, as opposed to our believing them to be true, would follow from that. Since we do believe them, we will assert them to be true. But no fact of the world would have been shown to be a necessary condition of our believing things about the world. This is still compatible with our insisting that we must take certain things to be true of the world in order to see ourselves as believing or having any opinions or attitudes about anything.

The fact about interpretation and the ascription of belief is the important point. I think it is enough in itself to prevent us from ever arriving at the "subjectivist" picture of the world. By that I mean only that we could never consistently arrive at the "subjectivist" conception of values or color, not that that conception is false, or necessarily false, or a contradiction, as it would have to be on the stronger conclusion that sees a necessary connection between the body of our beliefs as a whole and their truth. That stronger conclusion would say that there could not possibly be such things as evaluative beliefs, or color beliefs, unless they were

on the whole in fact true. And that is the kind of necessary connection that only idealism or some form of “subjectivism” would seem able to explain.

To think that we are forced to that strong conclusion by what I am calling the fact about interpretation would be to take the kind of step Bishop Berkeley took to his form of idealism. Because he thought that we cannot conceive of an object without perceiving it, he thought that we cannot conceive of an object that remains unperceived. He concluded that an object could not possibly exist unperceived—it is inconceivable. But that is to start with the fact that we cannot do something, that we cannot perform a certain feat, and to conclude that a certain thing could not possibly be so. If we distinguish, as we must, between what we cannot do and what cannot be so, between what we cannot consistently think and what is in itself inconsistent, the fact about the conditions of interpretation will not support the idea that our color beliefs or our evaluative beliefs are simply such that most of them must be true. That conclusion anyway would once again encourage idealism or “subjectivism.” But the fact about interpretation would mean that that is no threat. We would be in no position to deny or refute “subjectivism” on the grounds that it is inconsistent, but we would never be able consistently to reach the thought of the truth of “subjectivism” either. Denying “subjectivism” is not the only way of avoiding it.

So although I strongly resist the “subjectivity” of value, I do not wish to be understood as defending the idea that on the contrary values are “objective.” The tendency to draw that conclusion directly from the unacceptability of “subjectivism” is good evidence, if more were needed, of the power of that traditional metaphysical dichotomy. It still represents the structure in terms of which we want to understand things. We feel that colors, or values, or whatever it might be, must be either “subjective” or “objective.” We think that either there is something corresponding to our thoughts about them in the fully “objective” world or

there is not. So if we cannot really understand values to be “subjective,” we think they must be “objective” after all.

I believe that we cannot get a satisfactory understanding of ourselves in that way. That leaves us dissatisfied. It seems as if it couldn't be simply impossible. So we persist. And once again we apply the traditional dichotomy. This tendency, I believe, is the place to look for the real source of “subjectivism” about values. It would help explain why we can expect that some form of the view will always be with us.