

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

DOES MORAL PHILOSOPHY REST ON A MISTAKE?

W.T. Jones



HUMANITIES WORKING PAPER 132

March 1988

DOES MORAL PHILOSOPHY REST ON A MISTAKE?

W. T. Jones

California Institute of Technology
Pasadena, CA 91125

I have expropriated the title of Prichard's 1912 paper¹ because, I too answer his question affirmatively. But the mistake I detect is not the one Prichard thought he had uncovered, and his article is a classic example of the mistake I propose to discuss. It is to believe, as some moral philosophers still appear to do, that moral philosophy has a special domain or special method that distinguishes it in some important way from sociology, anthropology, psychology and economics. I shall argue that these moral philosophers are misled by the "philosophical" vocabulary they use.

I have two reasons for recommending to moral philosophers that they abandon such terms as "moral principle," "moral law," "moral intuition," "original position" and the rest of the traditional vocabulary. First, it encourages them to sublime—to etherealize—their discipline—by constructing over-arching, but empty, theories, and so distracts them from developing the kind of small-scale theory that would help generate the guidelines needed by people in moral quandaries.

Second, the moral-principle vocabulary misleads them. Rather than referring, as they believe it does, to features of the world that are distinctive in some way, different from the features of the world referred to by such social-scientific terms as "social norm," "decision rule," "perception," "starting point for negotiation," it only expresses attitudes of the moral philosophers—specifically, the affective states aroused in them when they think about the difficult decisions that they and other people must make in this uncertain world.

Thus this paper is addressed only to moral philosophers, and indeed to them only when they are doing moral philosophy.² I am not, that is to say, recommending wholesale elimination of the traditional language. Most people who are concerned about doing the right thing need some psychic support when they are at difficult decision points; at such times the traditional language is likely to be helpful. Moreover, a complex web of social, political and legal institutions, not to mention literature and the arts, are embedded in and presuppose the traditional language. A linguistic revolution that swept all this away, if *per impossible* it could be carried out, would have an enormous and unpredictable impact on the culture. In contrast, the revolution I am proposing is very modest; it would affect only the way in which moral philosophers talk moral philosophy with each other.

Thus the thesis of this paper is that the discussions of moral philosophers amongst themselves would gain much and lose nothing of consequence if they abandoned the traditional vocabulary. After setting it out (Section I), I shall illustrate it by giving sketches, in a social-roles language, of two features of the moral life that moral philosophers usually discuss in a moral-principles vocabulary: a suggestion about the nature of moral commitment, a suggestion, that is,

about why people believe they have an obligation to do so-and-so (Section II) and a suggestion about the causes of moral change, a suggestion, that is, about why, on occasion, they cease to believe they have an obligation to do so-and-so and come to believe they have an obligation to do such-and-such (Section III).

I shall then try to anticipate what I believe are the main criticisms that moral philosophers will level against this reduction. This discussion falls into two sections which reflect the differing reactions of moral philosophers to what Alastair MacIntyre calls "the Enlightenment Project."³ He argues that the project was a complete failure: it deprived moral rules of their former justification without being able to generate an alternative justification that has won general acceptance. Generally speaking, moral philosophers have reacted to this situation in one or the other of two ways. One group, whom I shall call realists⁴ and who for the most part are foundationists, either reject MacIntyre's conclusion that the Enlightenment Project has failed or, like MacIntyre himself, believe that another strategy of justification is possible. The issue between myself and these philosophers is of the deep kind that I call a difference in world view⁵ (Section IV).

The second group, whom I shall call nonrealists and among whom I include the constructivists, have abandoned attempts to justify moral rules and turned their attention to other problems, for instance, to uncovering the logic of those rules. My approach to this second group of philosophers will be to ask why, given their seemingly very different conception of the task of moral philosophy, they continue to use a vocabulary that is so closely associated with the realists' efforts to justify moral rules (Section V).

I

I start from the (obvious) fact that people assess their own and other people's actions both retrospectively and prospectively. They say such things as "that would be frightfully nice," "that was callous," "that would be indecent," "that was splendid," "that would be monstrous." Since the list of terms with which people praise and blame, point with pride and view with alarm, seek to encourage and to deter is enormously long, I shall concentrate on the terms in which people declare that they and others have a duty to do something or a duty to abstain from doing something. I choose this set of terms because moral obligation is a central concept for consequentialists and for neo-Kantians alike and because if the reduction of a moral-principles language to a social-norms language can be carried out for these terms, it should be possible to make a similar analysis of evaluative terms less saturated in the traditional platonizing world view.

I shall argue that the duties people feel, and the duties that moral philosophers tell them they ought to feel, to perform such-and-such acts and to abstain from such-and-such other acts do not arise from the nature of those acts, as the moral philosophers agree in maintaining (even whilst differing amongst themselves as to whether acts are duties because they involve treating other people as persons or because these acts produce the greatest good for the greatest number). For *any* act whatever that is expected by the members of some social group can come to be felt as a duty by the members of that social group, and any act that is frequently repeated in the appropriate circumstances will come to be expected by the members of the social group.

But though all those acts that are held to be duties, fulfill expectations, not all acts that fulfill expectations are held to be duties. What, then, are the differentia of the latter? I shall define these in the reduced social-roles vocabulary, and I shall start with a simple example. Consider a new dean, a young outsider, who has come into office after the long reign of an autocratic predecessor. Uncertain what will be acceptable to the faculty, he seeks the advice of department heads regarding an administrative appointment he must make. This creates an expectation, however slight, on the part of the department heads that they will be consulted, and that their advice will be taken, regarding subsequent administrative appointments—something they would certainly never have expected from his predecessor.

If he does consult them again this strengthens the department heads' expectations that the group will continue to be consulted. In due course, what was initially experienced merely as an expectation of future decanal behavior comes to be experienced as decanal duty—not merely something that he happens to do but something that he ought to do.

If the dean has internalized this expectation he will feel the same weight of censure (self-blame), whenever he fails to consult the department heads, that the department heads feel against him. Social censure and self-blame operate powerfully to induce the dean to satisfy, and in satisfying to strengthen, the department heads' expectations, and eventually the practice of consultation will be incorporated in the by-laws of the faculty and approved by the trustees. It is practice, then, and the expectations arising from practice, that make it binding on the dean to consult the department heads.

Thus a defeated expectation can generate moral indignation, but clearly, not every defeated expectation does. If the dean regularly arrives in his office at 9:00 a.m. people who have business with him will soon come to expect him to arrive at that hour and will plan their own day accordingly. If some day he does not turn up until noon they will doubtlessly be annoyed, but they will hardly feel indignation. What, then, are the differences in the social contexts of these two expectations that lead people to feel moral indignation when one of them is frustrated and only social annoyance when the other is frustrated?

II

This has brought me to the first of the two illustrative sketches I mentioned at the outset: I shall give an account, in a reduced, social-role language, of what people are feeling when they say, as on occasion they do say, that they are morally obliged to do such-and-such acts.

Note, first, that the affect that I shall call moral indignation is distinguishable from the affect that I shall call social indignation in much the same way that the affect called fear is distinguishable from the affect called anger. Though fear and anger are quite similar affects, they differ enough for one to be able to ask how the circumstances in which a person feels fear differ from those in which that person feels anger. One might form the hypothesis that people tend to experience fear in situations in which they feel insecure and anger in situations in which they do not, and it would be possible to check this conjecture by asking people whether they have sometimes found their anger changing into fear as they realize that the person who aroused their anger is in a position to harm them.

I shall adopt an analogous procedure with respect to the distinction between indignation that is moral and indignation that is social, and I shall suggest two features of the context that cause people to experience moral indignation in distinction from social indignation. Further, I shall argue that, just as it is possible to distinguish between fear and anger without assuming that one of those affects has a superior status (cosmically? metaphysically?) so it is possible to distinguish between social and moral norms without having to buy into moral realism, which, as I shall argue in Section IV, would involve having to swallow some very unpalatable beliefs.

What, then, are the two contextual features that result in an experience of moral, rather than social, obligation? The first is that people see violation of the norm in question as not merely a solecism, but as a threat to themselves or to what they hold dear. An example might be the hullabaloo fifty years ago over King Edward VIII's desire to marry a divorcée. Some Britons probably could not have cared less; others—amongst them possibly members of the Labour Party—were delighted at this sign of disrespect for the Establishment. Still others, however, elderly members of the upper class, for instance, experienced moral indignation at the violation of a norm that they took to be symbolic of their whole way of life.

Another example is the change in attitude toward smoking that is now occurring. Twenty years or less ago smoking in public places, so far from violating any social norm, conformed to one and was consequentially admired. More recently, and as a result, of course, of the increasing evidence of a correlation between smoking and lung cancer, smoking has come to be censured because it conflicts with a new social norm. Indeed, so widespread is the indignation over the adverse effects of smoking on the lungs of nonsmokers that it is moving rapidly toward becoming the violation of a moral, not merely social, norm. What, a short time ago, was social annoyance is becoming moral indignation. Thus norms are context-relative. What is a social norm at t_1 for I_1 may be a moral norm for her at t_2 and already a moral norm for I_2 at t_1 .

So much for the first feature of social situations that may cause what would otherwise be experienced as social indignation to be experienced as moral indignation—may, that is, cause what would otherwise be a social norm to be a moral norm. In order to understand the second feature, notice that when we are dealing with another person rather than with inanimate nature, we believe that her behavior *means* something, and also that it means to her pretty much what it means to us. If we didn't believe this we would never seek to interact with others; we would merely react to their behavior, as we react to a stone that is falling in our direction. In the case of people, as in the case of stones, we of course observe behavior, but in the case of people, in contrast to the case of stones, we observe behavior in order to be able to infer intentions, and we respond, not only to the behavior, but to what we take those intentions to be.

For instance, someone extends his right arm from his side; I read this as an intention to shake hands. Instead, his arm continues to move up towards his face—I see that his intention is to remove his glasses. I am mildly embarrassed, possibly disappointed, but I am not indignant. Moral indignation is one's response when the misreading involves a particular kind of intention, the intention to cooperate. Generally speaking, people feel moral indignation when (1) they have read a bit of someone else's behavior as meaning an intention to cooperate, (2) that person's cooperation would help them and his failure to cooperate would harm them, (3) it turns out they have misread it—it was not an intention to cooperate, and (4) they believe their original reading was justified—the

other person constructed his behavior carelessly, or he changed his mind without giving warning, or he wanted it to be misread.

When it is a matter of inferring, or reading, the intentions of people involved in any sort of joint activity—Hume's example of two men, each with an oar, engaged in rowing a boat is a good one—each party is likely to read the other's behavior as showing an intention to continue the activity that is under way. The possibility of social life depends on most people developing this kind of expectation in most situations involving joint activity; clearly, cooperation cannot wait on the successful negotiation of explicit contracts.

Three further comments: First, there may be disagreement among the participants about whether the pattern of behavior in which they are involved is, or is not, a pattern of cooperation. The department heads, believing that the dean's behavior in consulting them fits the pattern of cooperation, are indignant when he ceases to consult. He, however, may see his behavior as fitting the learning-to-do-my-job pattern, in which case his behavior does not in the least commit him to consult once he has gained the experience he thinks he needs. Here again, then, what is a social norm for I_1 at t_1 may be a moral norm for I_2 at t_1 and, quite possibly, a moral norm for I_1 at t_2 .

In cases of this kind much depends on whether one party to the disagreement can persuade the other that his understanding of the other's behavior was warranted. The critical question for each party is thus the other party's intentions: Did he mislead me, or did I misread him? If the two parties cannot agree about whether it is a case of misleading or of misreading their dispute may have to be referred to the courts, which, over the centuries, have developed rules for deciding whether such-and-such a verbal expression means that an understanding exists.

Second, explicit promises and formal contracts (the verbal expressions on which courts focus attention) do not underlie cooperation and "justify" it. On the contrary, behaviors that reflect and depend on social expectations can be seen operating in societies far too simple even to have conceived explicitly the notions of promise-keeping and contracting, let alone to have formulated explicit rules. But of course, once promise-keeping and contracting emerge as social institutions they provide alternative routes to commitment. And because people on the whole keep promises and fulfill contracts, violations occasion moral indignation, just as does any failure to continue cooperation once it has begun.

Third, cooperation does not actually have to be underway for one party to expect that the other means to cooperate. Suppose a patient fails to keep his first appointment with a physician. Though the physician has no reason, based on prior experience with this patient, to expect him to be punctual, the physician knows—and the patient presumably knows—that the role of physician and the role of patient involve cooperation at many points, including the making and keeping of appointments. It is enough if the actual behavior (telephoning, asking for an appointment, accepting it) fits the pattern of cooperation.

Thus, while I of course agree that people commonly distinguish between moral and social norms, this distinction can be captured without attributing any special transcendental status to moral norms, that supposedly distinguishes them from the social norms that moral philosophers readily admit vary across cultures. The moral necessity which people experience in connection with moral norms, is not in the norms, which are just whatever they happen to be, but in people's attitudes toward them. It is, as Hume wrote of natural necessity, something "which we *feel* in the mind,"⁶

specifically what we feel whenever we come to realize that behavior we had read as a sign of a readiness to cooperate was not such a sign.

To summarize: I have given an account of the difference between moral and social norms that parallels the analysis a psychologist might give of the difference between fear and anger, and has no more of a hidden metaphysical agenda. She might indeed suggest that anger is a "better" (because healthier or otherwise more useful) affect than fear, but it surely would not enter her head to argue that for this reason anger has a superior ontological status. Why, when they are discussing the difference between social and moral norms, do moral philosophers feel the need to make this sort of move? I shall return to this question in Section V.

III

I have now given a sketch, in a reduced, social-role language, of why people come to feel moral commitments—an account, that is, of why people believe themselves and other people to be obliged to perform or to abstain from performing such-and-such acts. I shall now give a sketch of how and why moral commitments change, the second feature of the moral life which I mentioned at the start of this paper.

The history of culture is strewn with commitments that have been abandoned, and not only abandoned, but repudiated: burning heretics, for instance, slaughtering Jews, drowning witches. How can all this wreckage be explained in the social-role language that I am recommending? Doesn't an explanation of moral commitment in terms of social roles lock the commitments, once made, into place?

This objection ignores the fact that a role is not a straightjacket; it is a loosely fitting garment, which the wearers are constantly adjusting, opening up a seam here, taking one in there. Adjustments of this kind—modulations, I shall call them—are intrinsic features of all roles, not occasional occurrences in otherwise fixed and static roles. Accordingly, obligation, understood in the reduced way I have proposed, so far from being an all or none affair, is a matter of differential weights that are attached to the norms that are embedded in, and that, collectively, form the role; and these differential weights vary over time and from role-player to role-player.

A modulation is large if it affects a norm that people think is important; hence large modulations usually raise what are perceived as "moral" issues, whereas proposed small modulations raise only "social" issues. Modulations, large or small, are initiated when a few members of some social group notice what seem to them to be relevant similarities and differences which most members of that group, having less sensitive antennae, have overlooked. And they are successful, if they are successful, because when they are called to the attention of the less sensitive, they are acknowledged to be relevant.

Thus, when Mr. Knightley pointed out to Emma the disparity between the role of lady and her behavior to Miss Bates—"I cannot see you acting wrong, without a remonstrance. How could you be so unfeeling to Miss Bates? How could you be so insolent in your wit to a woman of her character, age, and situation?"—Emma "was most forcibly struck. The truth of his representation there was no denying. She felt it at her heart."⁷ She saw that, to be consistent, she had to extend the scope of the norm "lady" to Miss Bates.

Suppose that the other members of Highgate society had also been cold-shouldering Miss Bates. It is possible that, like Emma, they would all see that Mr. Knightley's three similarities were relevant. If so, the scope of the norm "lady" would have been enlarged in that Sussex village. But had their view of "lady" differed markedly from Mr. Knightley's and Emma's, they might not have followed his lead. In that event, and despite his success with Emma, Mr. Knightley's perception of similarities would have been an aberration. For, though criteria of relevance obviously change over time, the criteria, whatever they are, that are perceived as relevant by the members of some social group *are* relevant for that social group at that time.

This applies as much to the perception of similarities (and dissimilarities) that are relevant to moral norms as it does to similarities (and dissimilarities) that are relevant to social norms. This is why I have brought Mr. Knightley, Emma, and Miss Bates into the discussion. Today that particular modulation, and the differentiations in the environment to which it was responsive will be viewed as trivial, though they were weighty for the social class of which Mr. Knightley, Emma, and Miss Bates were members.

The fact that what was once weighty can become trivial reinforces my thesis that there is no difference between the ways in which small-scale and large-scale modulations are introduced, and no difference in the ways in which they become established, if indeed they become established. Modulations that take on, and so become a part of the culture are now defended as correct usage by all the moral conservatives who, earlier, had condemned them as dangerous deviations.

Large-scale modulators—Beecher and his fellow abolitionists, for example— are not fundamentally different from small-scale modulators like Mr. Knightley. Both call attention to the fact that people who are outside some norm as it is currently understood (slaves; Miss Bates) are similar in a respect that the modulators regard as relevant to those inside the norm (human beings; ladies). Both point out that exclusion from the norm harms those who are excluded. Both feel strongly about the harm and try to communicate their feeling to others, whose changed behavior, if they are influenced by the modulator, will change the norm in the desired direction. Those already included in the norm are more likely than not to regard themselves as harmed—loss of privileged status, position, etc.—if those excluded come to be included, and they too may feel very strongly about this.

Thus the operative norm⁸ at any particular time may be thought of as a mean between the position of those who want to make the norm more inclusive and those who want to make it less inclusive—the "pro-life" lobby is a current example of the latter desire. Other examples are the norms that specify who is eligible to vote (whites only? also Blacks? men only? also women?), what a human being is (all fetuses? if not, how mature?), and the rights of animals (unconditioned? if not, how conditioned? all animals? if not, which? rattlesnakes as well as cats and dogs?). As of now, there is no consensus on these and many other issues. Public opinion will doubtless settle down, as proposed modulations are either accepted or rejected. But at least in the kind of society we live in today, no norm stays fixed for very long.

In a word, the mean moves over time, sometimes becoming more inclusive, sometimes becoming less. Once a substantial number of members of the social group have accepted the modulation and adjusted their behavior accordingly, it may begin to move rapidly. This can be accounted for in terms, once again, of social expectations. As one finds that more and more of one's

fellows expect one to accept the modulation, it becomes increasingly likely that one will find reasons for going along. Since most people carry around with them a reservoir of reasons,⁹ there is seldom much of a problem about finding them. Thus it is not difficult to account for the emergence of a consensus once a mean begins to move.

But why does the mean begin to move? Why do social expectations change enough for most people to want to go along? These are questions that, as it seems to me, moral philosophers should consider, for it is idle to ask what one's duty is with respect to x until a consensus of expectations concerning the performance of x has emerged. Philosophers ask questions about what one's duty is because they seldom take the fact of social change very seriously. And for the same reason their answers, however disguised they may be by the traditional timeless vocabulary, are usually in terms either of some earlier or of some projected future consensus of expectations.¹⁰

To summarize: When in the reduced language I am recommending, one equates an obligation to do \underline{x} with the social expectations for some role, one must bear in mind that this role is constantly changing as people respond in various ways, and at various rates, to the various modulations proposed by various modulators. If one slips into thinking of some role synchronically and looks at "instantaneous" slices of time, which is how moral philosophers usually look at roles, then the role looks stable and uniform; it changes, if at all, only between successive synchronic slices. It seems highly paradoxical to define obligations, as I do, in terms of social expectations and also to insist, as I do, that expectations change.

However, if one thinks diachronically, as social psychologists and cultural anthropologists tend to do, the paradox disappears. All sorts of modulations, large and small, are emerging, gaining adherents, losing them, gaining them again. . . . There is now no longer anything like an agreed-on set of expectations that completely determine every role player's duty. Instead, at any time a number of different sets of expectations are making claims against the role-players, some of whom are more responsive, some less responsive, some responsive in this respect, some in that respect, to the various claims. From this point of view the interesting questions are why, in such fluid situations, a consensus, i.e., a relatively stable obligation sometimes emerges; why it has coalesced on this, rather than some other, modulation; why it eventually dissolves. I shall discuss these questions briefly in the next section.

Meanwhile, what is the appropriate stance for moral philosophers to adopt, given the fact that norms continuously modulate? Since the traditional language powerfully suggests that norms ("real" norms) do not modulate, one could hope that they would abandon the traditional language. My two illustrative sketches have been written with this recommendation in mind. Freed from the traditional language, moral philosophers might move away from their fixation on high-level, very abstract theories of doubtful relevance to real life and turn their attention to devising lower-level guidelines—alternatively, "hypothetical imperatives"¹¹—that would be relevant to the problems of people in quandaries.

Even philosophers who are far less skeptical than I am of the possibility of bridging the gap between such lower-level principles and the super-principles with which moral philosophers usually busy themselves, might at least agree that small-scale theory, if not a substitute for large-scale theory construction, would be a useful propaedeutic.

IV

I shall now set out and reply to what I believe are the two chief objections that will be raised against this reduced view by the moral realists, as I have called philosophers who reject MacIntyre's Enlightenment Project root and branch.

First, they will argue that the reduced view cannot account for the categorical nature of moral rules. People, they will point out, want to justify their actions; they want good reasons for what they do. Unless there are moral principles on which the pursuit of good reasons terminate, the whole structure of reasons collapse. The possibility of justifying one's action, in distinction from merely finding support for it in public opinion, vanishes.

Second, they will argue that people have an obligation to perform certain acts even if these acts are not elements in any social role whatever, indeed, even if these acts are directly antithetical to acts which important social roles require. And they will maintain that this independent-of-roles feature of such obligations cannot be accounted for unless there are transcendental norms, i.e., norms that are independent of people's beliefs and expectations.

(1) My reply to the first of these two arguments is that the reduced view can provide as much justification for moral principles as they need and—as it were—deserve. Most people, certainly, want good reasons for what they do. The culture is a fertile supplier of reasons, and if some reasons have staying power it is because they appear to at least some people to be good ones. Most people, most of the time, can find several reasons for doing what they want to do, and they offer one or another of these reasons to different challengers, choosing for each challenger a reason that they believe will appeal to her as a good one. That is to say, the problem for most people most of the time is less to find a reason that satisfies oneself than to find one that satisfies other people, one that, above all, does not expose one to the charge of inconsistency.¹² People do not like to be told, "Yesterday you abstained from doing so-and-so. Today in an exactly similar situation you are doing so-and-so." Hence they will try to find a good reason for treating the two cases as "fundamentally" different, despite the "superficial" similarity which the objector has pointed out. If these reasons are challenged, further reasons can be produced.

Sooner or later, of course, the process of justification, the process of finding good reasons, like all processes, has to end—no further reasons can be adduced. At this point the challenger is either satisfied or she is not. Given the fact that the culture is a fertile supplier of reasons—that it is, simply reflects the variety of human nature and the complexity of the situations in which people act—it is unlikely that all challengers will be satisfied. If they are not there's the end on it. Moral realism does not bring it about that agreements are actually reached. The reduced view does not bring it about that agreements are not reached. Both views are compatible with there being a distinction between those who pigheadedly refuse to give any reasons at all and those who, because they look for reasons, sometimes reach an agreement.

People sometimes cut short the process of finding good reasons for a decision by appealing directly to a moral principle. I suggest that a moral principle is simply an expectation that has been abstracted from all of the roles in which it is actually embedded. Behaviors that are specific to one and only one role are not likely to be abstracted from the social role of which they are a part and elevated to the status of a principle, but behaviors that are expected in many different roles are easily abstracted and so come to be considered as "universally" or "categorically" binding, merely because

they are expected in many social situations and are so given a moral priority over any and all of the roles in which these expectations are actually embedded.

But when people to whom this view may seem plausible as long as it is stated in general terms are at choice-points in everyday circumstances (in contrast to being confronted with some contrived and artificial "case" that moral philosophers have concocted) they seldom look to moral principles for guidance. Rather, they ask themselves what is required by the specific social role—e.g., that of a father or that of a university president—they happen to be performing at the time they have to choose, for instance, between lying and telling the truth.

A university president who on occasion lies to his faculty without a qualm might not for a moment consider lying to his wife. A physician who would tell the truth about a very gloomy prognosis to a young man with family responsibilities (his good reason might be that the young man needs to get his affairs in order) might lie to an elderly patient with the same prognosis ("Let him die in peace"). And so on.

Thus when a norm ("Do not lie") is contemplated from a distance it may very well look universal, categorical, role-transcendent. But as one moves close up in order to decide how to behave in this or that particular situation, it turns out to have a series of local habitations and to be as context-relative as any so-called social norm. But the expectations from several similar roles in which one has performed in the past and in which one expects to perform again in the future accumulate and add weight to the expectations associated with the role in which one must act in the present. Violation of such a norm is experienced as a graver offense, and so as the violation of a moral, in distinction from a social, norm.

Few contemporary moral realists, I suppose, hold moral principles to be simple unqualified commands—"Never lie," "Never take a life." On the contrary, they are ingenious in working out the ways in which principles are qualified. But they appear to believe (1) that all the qualifications can be spelled out, at least "in principle," (2) that these "spellings out" are not exhaustively accounted for by differences in social roles, and (3) and that when they have been spelled out they are categorically binding on all people everywhere who find themselves in the spelled-out situations.

These assumptions are highly questionable. In the first place, what does the phrase "in principle" mean? How does it function in the kinds of sentences into which philosophers insert it? "Such-and-such is possible in principle" seldom claims that such-and-such is do-able if one has enough time. More often, it is a way of covering a retreat from an overextended commitment, a way of acknowledging, without explicitly admitting, that such-and-such is not do-able. In the second place, social practices, and the roles they generate, multiply so rapidly that, no matter how many qualifications moral realists introduce, they will never be able to take account of all the relevant differences across cultures, age cohorts, and economic classes.

The moral realists' examples are drawn from the opinions and practices of upper-middle class, late twentieth century Americans. But even with respect to such an unrepresentative subset of human kind, they seldom make any serious attempt to ascertain whether there is a consensus on any of the formulas that these philosophers hold to be universally recognized moral principles (for instance, "Lying is wrong, except the following circumstances. . ." "Life is sacred, except . . .").¹³

And verbal consensus on norms, were the moral realists actually able to show there is verbal consensus, would be no more reliable in revealing actual opinions and in predicting behavior than is

the virtual consensus of the affirmative answers to a pollster's question, "Do you want peace?" (or "lower taxes?" or "a balanced budget?"), as long as the answerers never consider what they might have to sacrifice for peace, lower taxes or a balanced budget. Thus I would say that Alan Schwartz's comment on the "Responsibility Principle" recently proposed by Talbot Page as a guide for judges in determining tort liability, applies generally—"incontrovertible on the level of abstraction on which [it is] stated, but the moral consensus . . . unravels in the face of actual cases."¹⁴

In a word, moral realists tend to be a bit naive methodologically. But that is not my present point. Rather, it is that one must distinguish between those moral realists who need not and those who should take account of actual practices and opinions, that is between moral realists who define moral principles as unqualified (without ellipses) categorical commands—a view associated, perhaps mistakenly, with the name of Kant—and those who define moral principles as qualified (with ellipses) categorical commands. The latter, so far as they actually begin filling in the ellipses with conditions and qualifications, make principles look less and less like moral laws having a transcendental authority and more and more like empirical decision rules that can be adequately discussed in the reduced language—rules that, by eliminating the need to make a decision *ab novo* at every new choice point, simplify life and reduce the amount of muddle, confusion and contradiction in the world. That, surely, is justification enough for moral principles, without imposing on them the impossible burden of qualifying as "ultimate" good reasons.

(2) I turn now to the second argument of the moral realists. In order to meet it I must give an account in the reduced language of the widespread feeling that people have duties (e.g., to free slaves) whether they, or indeed anybody at all, believe that freeing slaves is a duty.

To begin with, it is an oversimplification to think of there being a sharp dichotomy between a monolithic block of modulators advocating some change and a monolithic block of opponents of the proposed change. On the contrary, there are a number of people who are, to a greater or lesser degree, modulating; there are other people who are, to a greater or lesser degree, resisting the modulation; there are still others who are marking time, hoping that the modulation will either succeed or fail before they have to act. Further—a very important point—these differing attitudes toward specific proposed modulations always occur against an extensive background of agreement on other social roles.

For these reasons it is not unreasonable of modulators to expect at least some of their fellow citizens to adopt their modulation and hence to regard them as having a duty to do so. But for the same reasons it is not unreasonable of the modulators' fellow citizens to hold that the modulators have misread their intentions. Thus, to take the issue of slavery as an example, it was not wholly unreasonable of the abolitionists to have expected slaveholders to free their slaves. After all, the slaveholders were not a monolithic block, some were freeing their slaves in their wills; more felt uneasy about the institution of slavery, fearing that it might be inconsistent with their beliefs about civil rights. The abolitionists saw all this but did not see that though the slaveholders valued consistency, most of them valued their property even more. Hence, when the slaveholders opposed emancipation, when the abolitionists' expectations were defeated, it is not surprising that they experienced moral indignation.

Today, when people—ordinary citizens as well as moralists—argue that freedom and equality are transcendental rights, i.e., that the slaveholders had a transcendental duty to free their

slaves and that they had this duty before they or anyone else believed they had it, these people are simply resonating with the abolitionists and sharing their disappointed expectations. Freedom looks to them like a transcendental norm because, knowing themselves to be living in a later time period than the one occupied by the abolitionists and the slaveholders, they believe themselves to be evaluating slavery dispassionately, i.e., "objectively." But—this is my point—phenomenologically speaking, they are not outside the time period containing the abolitionists and the slaveholders. They are very much inside it, united in sentiment with the abolitionists and sharing their indignation over their defeated expectations. So far from having discovered timeless, transcendental norms, they are simply identifying with the operative norms of the abolitionists.

Thus an alternative account is possible of the widespread belief that duties (e.g. the duty to allow Blacks their civil rights), are role-independent. It will not appeal to the moral realists because there is a hidden personal agenda in their thesis that moral norms have a transcendental source. It is not so much that they are seeking, as they say they are, to justify the abolitionists' modulation—the enormous changes in attitudes and institutions that their initiative helped launch are surely justification enough. It is, rather, that they want to allay their anxiety about their own decisions at whatever choice points they may personally face. They can do this only if they feel assured that at every choice point there is *something* that is the right thing to do, even if they cannot always be sure what it is.

Dewey and, earlier, Nietzsche recognized that the impetus for much philosophical thinking is a quest for certainty. The moral realists are a case in point. They are made uneasy by moral disagreements that terminate only because the disputants get tired, or work out a compromise, or agree to disagree. They want it to be the case that at every choice point there is a course of action that is the right one to choose, so that, though the disputants may not hit on it, there is at least the possibility of moral disagreements terminating because all parties have come to see what the right action is.

The moral realists find the reduced view which I accept so easily, intolerable because they have a low threshold for living with moral dissonance and because on the reduced view it is not merely the case that it is often difficult to decide what one's duty is; rather, the *concept* of obligation is problematic. One can of course know, before a modulation begins to win adherents and after it has been widely accepted, what one's duty is. But when a consensus has been disrupted and before a new one has been achieved it is impossible to know what one's duty is, for at such times one's duty is indeterminate. At such times one can either (1) align oneself with some proposed modulation, hoping that one has guessed right and that it will prove to be the wave of the future, or (2) one can try to wait it out. Neither of these alternatives affords much comfort to people in quest of certainty.

We have reached the very deep difference in world view that I mentioned at the start of this paper. The moral realists see the world (this of course is the way I try to explain to myself what seems to me the extreme oddity of their view) through a platonizing lens; whatever the lens through which I perceive the world, it is not platonizing. To me, through my lens, it is as evident that there is no objective, transcendental moral realm as it is evident to them through theirs that there is one.

To the realists, I am trapped in a hermeneutical circle of beliefs, and beliefs about beliefs. Whether the beliefs are mine alone or those of my social group, or even those of all human beings (as might conceivably be the case) is immaterial, for without an outside reference point in something

which is not a belief, I am lost in subjectivism.

To me, a hermeneutical circle large enough to move around in is a comfortable place in which to engage in my traffic with nature and with other people. The realists are persuaded that belief systems that are not either suspended securely from the ceiling or resting firmly on the floor will collapse; I am persuaded that no belief system—large or small, chaotic or well organized—is ever permanently anchored; all, even those seemingly most firmly fixed, are sooner or later seen to have been detached, and floating, the whole time.¹⁵

The realists see a fundamental metaphysical similarity between nature and morals; I am struck by the difference. I agree of course that for a very long time there was a close parallelism between beliefs about nature and beliefs about morals. Up until the period in which what MacIntyre calls the Enlightenment Project was well underway there was no reason not to believe in the existence of a moral order that was as objective as, even more objective than, the natural order. The problem for the moral realists—and it is this, of course, that makes them so hostile to the Enlightenment Project—is the contrast between the way in which, during the past 500 years, beliefs about nature have tended increasingly to converge and the way in which beliefs about morals have remained diverse and conflicting. Since 1600 people have had a reason, which strikes them as a good reason, for concluding that there is a physical nature that is independent of our beliefs about it, a reason which they did not have earlier and which they still do not have for believing in the existence of a correspondingly independent moral realm. The parallelism that was "obvious" before 1600 has long since disappeared, at least for those who do not view the world through a platonizing lens.

Nor, finally, will there be agreement over whether the reduced view is dangerously subversive. I agree that a political elite—like the one Plato sketched in *The Republic*—would be a great danger if it were to adopt an esoteric language which could be used to manipulate the opinions and practices of the masses. But the chances of moral philosophers becoming a political elite are sufficiently remote for worry on that score to be a sign of neurosis, and there seems no more to be feared from interchanges by moral philosophers in a technical language than from interchanges by astrophysicists or neurophysiologists in the esoteric languages that they use to discuss the problems they encounter in their disciplines.

I take more seriously the argument that the reduced language I have proposed for moral philosophers might have adverse social consequences if it came to be widely adopted, for many, perhaps most, people are risk-averse, and uncertainty may induce indecision, weakness of will. Over the centuries, the authoritarian religions were a defense against scepticism. Now that religious faith has been weakened, the beliefs of the moral realists, offering an assurance that the moral realm is rational, are one of the few remaining bulwarks.

To this line of argument I reply that this article will not be widely read and that most of those who are made uneasy by it have at their disposal, as do the moral realists, defensive strategies which will protect them from my thesis. Certainly it is true that modulations succeed or fail in part because people become aware of empirical evidence that seems to them relevant—the changing attitude toward smoking is an obvious example. And the strain for consistency, to which I have already referred, is also a factor.¹⁶ But the larger the proposed modulation, the less likely it is that evidence and logic are determining. When it is a matter of deep, strongly cathected attitudes, people

are less moved by argument than by something like esthetic taste. The model for understanding profound modulations, such as the one being discussed here, is surely more like the shift from Mozart to Schumann or from Gerôme to Manet—than it is like the correcting of a faulty inference. To me this seems an adequate response to the charge of subversion. To the moral realists it will be very lame. Here again our arguments—theirs and mine—slip past each other without engaging or even making contact. Deep differences in world view generate non-terminating disagreements.

V

I turn now to a group of moral philosophers who operate (more or less easily) within the climate of opinion of which MacIntyre's Enlightenment Project was an early phase. Whereas the contention of the realists is, essentially, that I have not grasped what moral philosophy is about, the nonrealists, as I am calling the latter group, will say that they have long been doing just what I am urging them to do; I have been misled by their preference for using the traditional language to discuss what I want them to discuss.

Perhaps the difference between myself and the nonrealists is merely semantical. I hope so, but I doubt it. For, as I have argued elsewhere, differential preferences for one term or another, when two terms refer to the same object, always reflect differential attitudes toward the object named and that sometimes these different attitudes are so deep that they amount to differences in world views.¹⁷ In an effort, therefore, to discover how disturbingly deep, or how satisfyingly shallow, is the difference reflected in my preference for the reduced language and their preference for the traditional language, I shall ask why they persist in using the traditional language.

I shall begin by saying I believe retention of an otherwise outmoded vocabulary is usually a sign that this vocabulary is responsive to attitudes which are still entrenched in these users' *mentalité* despite their belief that they have abandoned them. With this as a working hypothesis, it is natural to ask what those entrenched attitudes are, and the answer is that they vary from case to case. For instance, one notices that men who believe that they have shed male chauvinism nevertheless cling to a vocabulary containing such words as "postman," "policeman," and "chairman." Although inertia may account for this usage in some cases, it is surely not implausible to suspect that at least some of these men are not yet fully liberated. Analogously, I shall argue that the traditional moral vocabulary survives in the writings of the nonrealists, who believe they have put foundationism and metaphysical ways of thinking behind them, because it enables them to enjoy metaphysical cake without having to eat it.

Consider, for a start, the non-realists' discussions of moral rules. One can read these discussions with an eye to whether they have got the structure of these rules "right" (which is the way they expect their discussions to be read and the way most philosophers in fact read them). Or one can read them, as I propose to,¹⁸ with an ear for the subtexts, or latent meanings that their discussions express, in the same way that a persistent use of masculine nouns expresses a latent hostility to the Women's Rights Movement. Read in this way, much that is puzzling about the non-realists' texts becomes transparently clear. They are animated by the same metaphysical anxieties that drive the realists to foundationism. The traditional language, functioning as a subtextual level, helps the non-realists allay these anxieties, whilst, at the textual level, nothing could be farther from their thought.

Where, for instance, from the point of view of the non-realists, "are" the moral principles discussed in their texts? Presumably, since the non-realists have abandoned transcendental norms, their principles are embedded in people's practices, and can be extracted from those practices inductively. Thus it is possible to "construct" the operative norms¹⁹ of a society from the actual practices²⁰ of the members of that society, as linguists construct linguistic principles by observing people's speech practices. But the non-realists do not proceed in this way. Instead of starting with practices drawn from the behavior of a variety of social and ethnic groups and then moving inductively to provisionally formulated norms, they start with norms in which *ex ante* they seem to be supremely confident and then proceed almost anecdotally, employing what can best be described as a strategy of selective exemplification. Moreover, they often rely on anecdotes about science-fictional creatures or on extremely artificial earth-bound cases to convince their readers.

This otherwise curious procedure can be explained if one assumes that, despite protestations that they have modulated moral philosophy away from foundationism and in the direction of the empirical relevance, the nonrealists' discussions of the structure of moral principles are animated by a vision, which they share with the nonrealists, of the domain of human decisionmaking and institution-design as through-and-through "rational", as systematically well ordered. In a word, they too sublime morals, and in a way that parallels the subliming of logic that Wittgenstein deplored. His comments on the latter apply to the former: to those in the grip of this conviction, moral philosophy, "seeks to see to the bottom of things . . ." They believe the moral realm "must be *utterly simple*," "*prior* to all experience," "of the purest crystal."²¹ Moral philosophers are subliming when, for instance, they take the strain for consistency, which operates differentially in different societies, as at once licensing their search for a single, all-inclusive and well-ordered system of moral rules and also as guaranteeing that they will find it. Thus the non-realists, as much as the realists if perhaps less obviously, are distracted from discussing the kinds of low-level guidelines that would be helpful to people in moral quandries.

In their tendency to sublime, moral philosophers of course are not alone. Some economists and political scientists also sublime, but whereas these scholars come out of a social-scientific culture which is to some extent a counter-weight to the vision, the nonrealists, like the realists, come out of a platonizing culture that was the natural habitat of abstract principles. There is nothing, therefore, to discourage the nonrealists from drifting to high-level generalizations and ignoring the messiness of actual practices, the muddle of real-life opinions.

Thus expressions like "moral principle" function in a more complex way in the texts of nonrealists than they do in the texts of realists. In the texts of the latter "moral principle" names whatever is named in my text by "transcendental norm"—if anything at all is named by the latter expression. In the texts of the nonrealists "moral principle" seems to refer to what, in my text, is called a decision-rule, or operative norm. But in their texts "moral principle" is also a part of a subtext, for at this point texts of Descartes and other early modern foundationists have been displaced—to borrow a term favored by deconstructionists—into the texts of the nonrealists.

I suggest, then, that the nonrealists' preference for the traditional vocabulary can be accounted for by the fact that it is saturated with the odor of realism. Whereas "operative norm" is associated with an empirical stance and so invites investigation of the possibility that in different societies, or even in the same society, different, and incompatible, norms may be operative, "moral

principle" is associated with a platonizing stance and so protects the nonrealists from having to face the possibility of there being irrational choice-points, choice-points at which good reasons are to be found for choosing any one of several alternatives but no better reason for²²

Again—to note another point at which the texts of moral philosophers carry a subtext—calling truth-telling and other operative norms "moral principles" makes them look like categorical imperatives. Inasmuch as in normative social science the goals—less crime in the street, a balanced budget, full employment, for instance—are all evidently hypothetical, this displacement gives moral philosophers a seeming warrant for distinguishing their discipline from the part of social science that deals with normative issues.

The cost of this declaration of independence is high. Calling truth-telling, for instance, a categorical imperative encourages the tendency of moral philosophers to indulge themselves in subliming and so leads them to formulate ends that are too abstract, grand, and remote to be helpful to people in moral quandaries. What is the use to a pregnant teenager, who is trying to decide whether to have an abortion, to be told to respect persons or, alternatively, to maximize happiness?

I shall conclude this paper with three more examples of subtexts that, as I believe, are to be found in the texts of nonrealists. First, there is the use of a floating (as I shall call it) first-person plural pronoun, which makes the ancient *consensus gentium* argument available in a contemporary context. This argument—which, like the teleological proof of the existence of God, has psychological appeal disproportionate to its evidentiary weight—was attractive, and presumably is still attractive, because a rule on which all mankind agrees seems to have an authority that cannot be traced to political institutions or historical traditions, which vary from tribe to tribe.

Nowadays an explicit appeal to a *consensus gentium* would convince few people. But the psychological appeal of the argument can be achieved indirectly, by a floating "we." Left unchallenged, it soars upward, in the texts of moral philosophers, from "you and I," to "our set," to "people like us," to "twentieth-century Americans," and so into the stratosphere of universal human nature. But if it is ever challenged, "we" can glide down to the level at which the writer believes the challenger may allow that the claim for a consensus is not wholly implausible. Since the moment the skeptic's back is turned, "we" unobtrusively floats up again, a *consensus gentium* is suggested without being asserted.

Much the same can be said about "intuition." Unlike the realists, who can consistently hold that people intuit transcendental norms, nonrealists, if queried, are likely to say that by an intuition they mean a feeling or perception, the off-the-top-of-the-head reaction of some ordinary chap. But, they continue, there is a principle embedded in the feeling which, when analyzed will prove to be embedded in the often superficially different intuitions of other ordinary chaps.

Taken at face value—that is, at the level of text—"intuition" thus makes only an empirical claim, i.e., a claim that further investigation will confirm or disconfirm. But "intuition" also functions in a subtext, for it is a term with a very long philosophical history, one which claims certainty for what is intuited and which holds intuitions to be the foundation on which knowledge (in this case, knowledge about how to live virtuously) rests. "Intuition" thus allows the nonrealists to imply, without explicitly saying, something they very much want to believe.

Finally, consider "original position." What does "original position" do for moral philosophers that a reduced phrase like "starting point for negotiating a change in relationship" cannot do? It straddles the chasm between the pre- and the post-Enlightenment *mentalités*. As with the floating "we," the meaning of "original position" can move from the empirical (in which case there may be quite a number of different and competing original positions, depending on which culture is the basis for one's "construction") to the transcendental (in which case there is one and only one *ur*-position). It can shift back and forth, as need arises, from the set of conditions that would exist if some particular society were to carry through a consistent application of its own basic assumptions to the set of conditions that generates a "really" just society and so holds for all societies everywhere.

The association of "original" with a pre-lapsidarian state of affairs is a subtext in which the idea of "original" not merely as first but as uncorrupted, gives the expression an ambiguity that appeals to nonrealist moral philosophers.

I conclude, then, that the traditional vocabulary survives in the writings of the non-realists because it enables them to make a transition, which is "imperceptible but . . . of the last consequence,"²³ from (1) "is believed by the members of some social group to be obligatory," which is an empirical claim to (2) "is obligatory," which is a nonempirical, categorical claim.

This transition is of the last consequence because, so long as it is imperceptible to the nonrealists, they can feel they have escaped from the hermeneutical circle which alarms them as much as it alarms the realists. But whither, into what third realm, beyond the physical realm and the realm of human experience, can they escape? Pending a satisfactory answer to this question—and the nonrealists, in my view, should tackle it directly, not bracket it as somebody else's problem—I shall continue to hold that moral philosophy rests on a mistake.

FOOTNOTES

- * I am much indebted to the following friends and colleagues for comments on earlier drafts of this paper: Nancy Atkinson, Jay Atlas, Brian Barry, John Benton, Sharon Bishop, Bruce E. Cain, Randall Curren, Alan Donagan, Burton H. Klein, Arthur Kuflik, Stephen J. Morse, Alan Schwartz, Eleanor M. Searle, Robert M. Stewart, Peter Westen, James F. Woodward, and Charles Young.
1. H. A. Prichard, "Does Moral Philosophy Rest on a Mistake?" *Mind*, N. S., vol. 21, 1912.
 2. See below, pp. [12-13, 21-22].
 3. I would not myself talk about a "Project" at all. "Project" seems to suggest a well-bounded, well-defined undertaking supervised by people who knew what they were about. That is an *ex post*, historian-type assessment of what, for insiders, was a whole complex of messy, often conflicting efforts, themselves only a phase of that still, more complex, and even more loosely bounded collection of family resemblances called the Enlightenment.
 4. I have no doubt that some, perhaps all, of the philosophers I assign to this pigeon-hole will dislike this label. Although I have consulted friends and colleagues I have not found a term which will not give offense in some quarters, and, as Sweeney pointed out apologetically to Doris and Dusty, "I've gotta use words when I talk to you."
 5. See, for instance, "Philosophical Disagreements and World Views," *Proceedings and Addresses of the American Philosophical Association*, vol. 43, (1970), "Reason, Feeling, and World View," in *The Philosophy of Brand Blanshard* (1980); also *Current Anthropology* and Jaynes pieces.
 6. *Enquiry*, Sec. 7, Pt. II.
 7. *Emma*, vol. III, ch. vii. If readers prefer a real-life example to a literary one, Mr. Justice Bazelon's address at the 1981 annual meeting of the American Psychological Association illustrates the way a modulation can start. Relying on his long experience hearing cases in which psychological "experts" testify, the judge discussed what he called the "sins" of the profession: a tendency to make "conclusory pronouncements," a failure to "expose the facts under their conclusions" and "the values underlying their choices of facts," and a failure to "come clean on the uncertainties of opinion that may exist. . . ." The fact that extensive excerpts from the address were printed in the *APA Monitor*, (vol. 12, no. 10) shows that this modulation is beginning to spread, though it is far from clear that it will eventually take on.

8. An "operative norm" is to be distinguished from an actual norm. The latter is a generalization about how people actually behave in such-and-such circumstances; the former is a generalization about how people should behave, i.e., a statement of a belief about the *expectations* of a social group in those circumstances. Since people's beliefs about how they and other people should behave affect how they behave, operative norms and actual norms tend to converge. But in some cases—in the area of sexual behavior, for instance—in some societies at some time they may diverge a great deal.
9. See below, [p. 14].
10. See below, [p. 19].
11. See below, p. [26].
12. There is in most, perhaps in all, societies what can be called a strain for consistency which takes the form of an inclination to treat similars similarly. But societies differ about the respects in which similars are similar, about how similars must be to be treated as similar, and about the way they rank-order consistency as a desirable.
13. The ellipses are intended to permit casuists to insert all the qualifications they wish to include before they declare that the principle in question is absolutely and categorically binding.
14. A. Schwartz, "Responsibility and Tort Liability," *Ethics*, vol. 97, no. 1, October 1986, p. 272.
15. Cite Piaget's metaphor in *Structuralism*.
16. See note 12.
17. Cite relevant discussion in TS&TH.
18. In "Philosophical Archaeology," (forthcoming in *Metaphilosophy*) I have sought to illustrate and justify this way of reading philosophical texts.
19. See note 8.
20. Practices like lying to A and telling the truth to B; opinions like, "Well, I almost never lie to A-ish people though I sometimes lie to B-ish people"—without formulating the A-ish and B-ish qualities she has in mind and without specifying what she means by "almost never" and "sometimes."

21. P.I. §§89-91, 97
22. Guido Calabresi's recognition of this accounts for the acute distress that suffuses the pages of *Tragic Choices*. For Calabresi, at least, the vision splendid has faded into the light of common day.
23. *Treatise*, Bk III, Pt. I, Sec. I.