

The Meaning of Norms and Values

By Chris Wright

Prefatory note from [my blog](#):

Back in my twenties, when I wrote this paper, I was consumed by philosophy, jumping from topic to topic in an effort to answer an implausibly wide range of questions. Unfortunately, my writing wasn't always "polished"; it was intuitive, dense, rough, difficult to read, not nearly as discursive and diffuse as that of most academics. So it was rarely publishable. The excerpts below from a long paper on values are an example. They're in the anti-"objectivist" spirit of [these simpler thoughts](#). I hadn't read J. L. Mackie's classic [Ethics: Inventing Right and Wrong](#), but, in retrospect, my arguments were broadly in agreement with his. There are no "objective truths" about values, about which values are right and which wrong, which good and which bad. People can have whatever values they want: they can believe Hitler and Stalin were *good*, and they aren't "wrong"—except in relation to other values, such as being compassionate, not killing millions of people, etc. There is no such thing as (literal) "evil," a remnant of theology: *objectively bad* is as meaningless as *objectively good*. We're dealing with subjective reactions here, subjective orientations toward a given object; the universal mode of experience according to which the value we assign to something *actually inheres objectively in it* is mistaken. Beauty, ugliness, goodness, magnificence, badness, and all other values are things we subjectively project into something; they aren't there independently, even though it may seem when you look at a beautiful face or painting that its beauty is objectively given, so to speak, just as objective as the measurements between features or the various colors that strike the eye. We think of the serial killer Jeffrey Dahmer as somehow *bad in his essence*; but in reality he was just someone who had a taste for killing and eating people, as others have a taste for gardening or growing their own vegetables or playing tennis or whatever. It is we who assign the negative value to Dahmer's activities, as we assign a positive value to, say, a medical doctor's activities. There would be nothing incoherent or objectively wrong about reversing the value-judgments: we're free to value whatever we want. It so happens, fortunately, that the vast majority of people share common intuitions about what is good and what is bad, so we all agree Dahmer was bad. But the only *genuine content* of our judgment is that *we don't like* his behavior. Our judgment has an "objective" form ("*He was...*"), but really we're only saying something about ourselves, not him.

In fact, the fundamentals of how we experience the world are mistaken. One might subsume them all under the category "objectivist fallacies." Our brain constructs our experience on the basis of a wealth of sensory information that it organizes according to innate and acquired schemata, principles/rules of perception, cognition, emotion, aesthetics, and so on. This is the central lesson of the brain sciences and psychology. "In itself" the world is an incomprehensible, colorless mess of elementary particles, electromagnetic waves, atoms and molecules built up into larger structures, sound waves, dark matter and dark energy, infinitely many things zipping around, colliding, attracting, etc. Out of this welter of confusion, the brain constructs our ordered

experience. (Evidently all our brains operate in virtually the same ways, since we all have virtually the exact same perceptual experiences. We're genetically almost identical to each other, with variations around the edges.) Our experience of the world is more immediately our experience of our own minds, the perceptual, cognitive, and emotional architecture of our minds. Thus, when you emotionally or aesthetically react to something—like when you react to a person in some context, your reaction is a function of your mood, your expectations, the impressions you've already formed from previous encounters, whether he or she reminds you of other people you've known, etc. There may be *some* validity to your judgment—for example, insofar as others would agree with you—but ultimately it reflects *you* as much as the other.

We interpret things as being *out there*, but in a sense they're really "*in here*." Here's a mundane example: sometimes if I'm reading a passage really fast I see a word as such-and-such, like, say, 'cautious.' But then I reread it and I see it's really 'caustic.' I remember having *actually seen* the word 'cautious,' but now that word has suddenly turned into 'caustic.' And I'm made vividly aware how completely the brain *constructs* everything I experience. I thought the word was *out there*, on the page, but it's really *in here*, in my mind—as, indeed, is the page itself, the page as I perceive it, along with all other objects. The world *as it is in itself* doesn't have a white page with black marks on it; it has only the entities that physics postulates. (As you can see, I'm a [scientific realist](#), a position that strikes me as common sense. See also [these reflections](#).)

Anyway, in the excerpts below I go into some detail about the nature of values and their lack of objectivity. I also discuss the nature of reason, in part because that topic, too, was obsessing me but also because I thought it bore on the question of values. Specifically, questions about practical reason, as opposed to theoretical reason, bear on values, on action, and on justifying or criticizing courses of action. When and how can we criticize someone for acting irrationally or holding irrational values/desires? What does that mean? In brief, I argue that a person's values can be irrational *only in relation to each other*; for example, it's irrational for me to value Donald Trump's presidency if I also value democracy, compassion, and the well-being of humanity. In that sense, practical irrationality is possible (contrary to what David Hume thought). A person can also demonstrate *theoretical* irrationality by denying a sound chain of reasoning, for instance, or deriving a conclusion through a faulty chain of reasoning. (Actually, I think this is also a type of practical irrationality, inasmuch as everyone fundamentally values, whether they know it or not, logic and the idea of being correct; so they're acting in contradiction to this more basic value by disregarding logic in a particular instance.)

There's a lot of lengthy, and doubtlessly excessive and repetitive, probing of various issues, such as the falsity of realism about values, the consequent semi-meaninglessness of value-judgments, the notion of "psychological consistency" (as distinguished from strict logical consistency) that I introduce to explicate practical irrationality, the nature of action, Kant's ethics and why they're wrong, and, at the end, certain implications of my arguments.

Just one more anticipatory note: the reason I say value-judgments are half-meaningless is that their 'objective,' 'representational' linguistic form contradicts their actual content, which is really the statement that the object elicits a certain reaction in the speaker. "The painting is

beautiful" isn't a wholly meaningful statement; more meaningful would be to say "I find the painting aesthetically pleasing" or something like that. The latter focuses attention on the main point of the utterance, namely the speaker's subjective response; the former statement pretends that something called "beauty," whatever that may be, is *in the painting*, an essential quality of it. The very word 'beauty,' like 'good,' 'bad,' and all other value words, isn't wholly meaningful, being an illusory reification or hypostatization of a subjective reaction. You can spend hours thinking about the word 'beauty,' reflecting on it, introspecting, intensely contemplating it, trying to get clear on exactly what it means, what it *is*—"because it's so abstract, you see!"—but you won't get very far. The problem isn't its abstractness; the problem is its contradictory fusion of subjectivity and objectivity—a state of mind reified into a supposedly objective thing, a thing whose essence it is to cause that particular 'admiring' state of mind. No such thing exists; the very idea of it is meaningless.

More generally, platonism is wrong, though a rather psychologically natural philosophy.

In short, we're all walking around utterly deceived about the true significance of an enormous range of our experiences and judgments. The very phenomenology of our experiences deceives us. The human mind was made to deceive itself. We ourselves are not what we think we are, as I've argued elsewhere; we are not substantial selves. The Buddhists are right about that. The external, mind-independent world isn't what we think, objectively full of colors and smells and tastes and sounds; it is the exotic world of physics and chemistry, not the world of our brain-constructed experience. It's as if Descartes' famous "demon" is systematically deceiving all of us about almost everything—except the demon is just the human brain, deceiving itself, so to speak.

I hope you find some of the following at least mildly thought-provoking—not to mention decipherable.

*

The concept of normativity presents a number of philosophical problems. For example, what is the relation between values and norms? Are there objective norms, objective values? What does that mean? *What is* normativity, and *what are* value-judgments? What do they really mean? Is there an important difference between norms of reason and other kinds of norms? What is the difference between practical and theoretical reason? Is it possible to be *motivated* by *reasons*? If so, how? If not, what are the implications? Are moral values part of the "fabric of the universe", or are they purely subjective? If they are, does that vitiate morality? —The list goes on and on. The sheer number of questions might seem daunting, quite aside from their intrinsic difficulty. Nevertheless, in this paper I'll try to answer many of them.

If that seems over-ambitious, I would respond that the answers to such questions are necessarily interrelated; any thinker, therefore, who wants to answer *one* question—with any

degree of thoroughness—must answer many. On the other hand, undertaking such a systematic endeavor in one paper means that the arguments will at times be sketchy.

1. Preliminaries

What kinds of norms are there? *Prima facie*, there are moral norms, aesthetic norms, social norms, norms of contemporary science, epistemic norms, norms of reason, and so on. It is worth noting that scientific and epistemic norms (and norms of social science, of philosophy, of psychology, etc.) are actually norms of reason in disguise. That is, they are supposed to be derivable from the basic rules of reason as applied to particular states of affairs and types of truth (be they psychological, philosophical, scientific or whatever). It seems, then, that there are but four major normative categories: moral, aesthetic, rational, and social, by which I mean all the myriad rules about how, in light of propriety, it is appropriate to act in certain situations.

Normative judgments are expressed through the words ‘ought’ and ‘should’. They are prescriptions. As such, they are closely related to values. One cannot prescribe something without thereby valuing something. Every prescription is, in fact, either a kind of restatement of a value-judgment¹ or based on the consideration that the prescribed thing is the means to an end (or *constitutive of* an end²). By ‘restatement of a value-judgment’ I mean something like ‘the “flip-side” of a value-judgment’. For instance, “One ought not to lie” is, in a sense, another way of saying “Being honest is better than lying”. By saying the former, you commit yourself to the latter; and by saying the latter, you commit yourself to the former.

Evidently the reason a particular prescription is associated with a particular value-judgment is that they both express approval (or disapproval) of the same object. In themselves—i.e., without adding a middle premise—they aren’t *logically* related, but the middle premise is supplied so automatically by speakers that statements like the two in my example are treated as if they logically entail one another. Hence, someone who asserts one and denies the other is seen as irrational, insensitive to logic. Strictly speaking, though, he is simply denying the premise that one ought to do whatever is good; and if he denies the premise, he is logically free to deny the

¹ In most contexts in this paper, by ‘value-judgment’ I mean a statement in the form of “So-and-so is bad [or good, etc.]”, not a statement like “I approve of such-and-such”.

² For example, “say that my end is outdoor exercise; here is an opportunity to go hiking, which is outdoor exercise; therefore I have reason to take this opportunity, not strictly speaking as a means to my end, but as a way of realizing it”. See Christine Korsgaard, “The Normativity of Instrumental Reason”, p. 215.

conclusion (in this case, that people shouldn't lie). It would, therefore, be better to say he is acting in a *psychologically* self-contradictory way, so to speak, rather than in a *logically* self-contradictory way. For he is being inconsistent in his attitude towards honesty: he approves of it and disapproves of it at once.

To take another example, suppose I exclaim, upon seeing Michelangelo's *David*, "This is beautiful!"—but then proceed to smash it into bits with a sledgehammer. You will probably call me irrational, not because I have done anything *logically* 'self-contradictory' but because my mental states are flagrantly inconsistent with each other. That this still strikes us as an example of irrationality (even though nothing *logically* inconsistent has been done) shows that practical rationality involves not only respecting logic and induction but also being *psychologically* consistent. If you approve of a given object, such as a sculpture, unless you act according to what are generally thought of as the psychological implications of your attitude of approval, you will justifiably be called irrational.

I'll return to these thoughts when I discuss practical reason in depth. For now, my point is, first of all, that particular value-judgments are tied to particular prescriptions (and vice versa) via the relation of psychological consistency. 'Prescriptive (or normative) implications' of a value-judgment are, first and foremost, psychological implications, though they can be made logical *per se* by supplying one or more missing premises. Secondly, one way of being irrational is by violating psychological consistency. Insofar, then, as someone doesn't adhere to the prescriptive implications of his value-judgment, he is being irrational. If I have as an end financial success but don't take the means to achieve my end, I am to that extent being irrational—because I am being psychologically self-contradictory.

It is misleading, and probably incorrect, to say (as some philosophers do) that by having an end, I have a *reason* to take the means to my end. To speak of a 'reason' here is superfluous and wrong. Rather, *just by virtue* of having some end, I implicitly posit 'prescriptions', or 'regulations', on my behavior. Or, to make the point in broader terms: insofar as I value something, I posit corresponding norms (which can be 'psychologically derived' from my value-judgment) to regulate my behavior. These norms are psychologically implicated in my value-judgment. If I flout them, I am being inconsistent with myself and hence irrational.

It seems, then, that we already have a preliminary answer to the question of what explains normative force—i.e., the question of what normativity means, of what 'ought' signifies and

where its force comes from. At least with regard to value-judgments that one makes oneself, one ought to respect their prescriptive implications just to the extent that one values *rationality*. Exactly what rationality is will be the subject of a later section, but it's clear that it involves, at the very least, basic self-consistency. One's attitudes, beliefs, acts, values and so on cannot be explicitly mutually self-contradictory—in either a logical or a psychological way³—if one is to be properly called 'rational'.

2. A digression into the philosophy of language

[Deleted.]

3. Value-judgments and reason

Given that a particular assertion attributes or denies value to something, what exactly does such an attribution or denial consist of? In other words, what is a value-judgment? It's useful to note that every value-judgment can be rephrased as "[The object] *has* [the value]". "The mountain is beautiful" = "The mountain has beauty"; "The theory is right" = "The theory has the property of truth". Every value-judgment ascribes a property to an object. "It is wrong to lie" means "Lying exhibits the property of moral wrongness". Value-judgments, therefore, cannot be merely 'disguised' emotive utterances like "Boo lying!" or "Yay for this theory!", as theorists like A. J. Ayer would have it.⁴ In fact, any kind of noncognitivism faces the difficulty of having to explain how it is that value-judgments can be so formally similar to scientific propositions while not aspiring to objective truth. *Prima facie*, both kinds of statements exhibit truth-conditions; both, therefore, are implicitly meant to describe truth, reality. They both '*represent*' a state of affairs, a state of affairs independent of the speaker's mere attitude or feeling. Lying *is* wrong, kindness *is* good. These statements can be rephrased using the word 'objectively' without any apparent violation of meaning: lying is (*objectively*) wrong. Noncognitivists may be right, say, that moral judgments involve accepting certain norms, or that they involve an expression of

³ As will be seen (and should be clear already), psychological consistency is the broader category. Logical consistency is merely the most obvious aspect of it.

⁴ See his book *Language, Truth, and Logic*.

a certain state of mind, but it is counterintuitive to deny that they also aim for objective truth in the way that scientific assertions do.

Perhaps that position is controversial. It is significant, however, that ordinary people act as if their value-judgments—especially their moral judgments—are objectively correct. People are as committed to their moral beliefs as they are to their scientific beliefs and their religious beliefs—that is, their beliefs about the way the world objectively is. They’re even committed to their *aesthetic* beliefs and values in that way. If you disagree with them that a certain song is beautiful they might well argue with you, try to convince you that you’re wrong, as a scientist might try to convince you of a particular hypothesis. In more philosophical moods they might say that judgments of beauty are merely subjective reactions with no claim to objective truth, but they almost invariably *act* as if their aesthetic judgments are objectively true. And as I said, this reaction is even more pronounced in the case of moral judgments.

In his paper “Objectivity and Truth: You’d Better Believe It”, Ronald Dworkin argues against moral skeptics who try to differentiate between two spheres of moral discourse, namely our ordinary, first-order, substantive judgments about which acts are right and which wrong, and the philosophical, second-order judgments about the claims to objectivity of the first-order judgments. Such skeptics deny that their arguments apply to the first-order judgments; their position is supposed to be neutral with regard to these. It’s only the meta-level philosophical claims they want to contest—claims involving phrases like ‘it’s really true’, ‘it’s true independently of us’, ‘it’s objectively true’. In short, these theorists profess to be skeptical only about the possibility of objective moral truth, and they maintain that such second-order positions are logically separable from first-order moral discourse (which they want to leave untouched). Dworkin disagrees, though. He thinks that second-order judgments are merely reaffirmations of the first-order ones, and that therefore the skeptics are inadvertently attacking the latter by attacking the former.

The most natural reading [of such claims as “My moral beliefs are really and objectively true”] shows them to be nothing but clarifying or emphatic or metaphorical restatements or elaborations of the proposition that [e.g.] abortion is wrong. If someone thinks abortion morally wrong, he might well say, for

example, in a heated moment, “It is just true that abortion is wrong”. But that would be only an impatient restatement of his substantive position....⁵

I mention Dworkin’s argument because it supports my own. Implicit in ordinary moral judgments is the positing of their objective truth.

Let’s assume, then, that value-judgments involve a claim to objectivity. To deny that would be to fly in the face of common intuitions as well as some persuasive arguments.⁶ In this respect, then, value-judgments are comparable to scientific hypotheses. So let’s look at the latter. Scientists believe that the most well-established hypotheses are universally valid, that they hold true for everyone because they hold true of reality itself. The truth of Darwinian evolution doesn’t depend on a person’s subjective reactions to it. Given its truth, scientists conclude that everyone *ought* to believe in it. Where does the ‘ought’ come from? Clearly by adding the further premise that people ought to believe true propositions. This premise, in turn, comes, via the mechanism of ‘psychological implication’, from the idea that it is good for one’s beliefs to be true. But at the moment, the relevant idea is that not only a scientist but *everyone* thinks that truth “lays claim to everyone’s assent”⁷—that it ought to be believed by everyone.

This is somewhat reminiscent of Kant’s idea that judgments of taste are universalizable, in that they make a claim on everyone. “In this respect, [they] are like ordinary statements of fact. Both are either correct or incorrect (‘valid’) independently of who makes the judgment (‘universally’).”⁸ Moreover, as Kenneth Rogerson points out, judgments of taste are, according to Kant, imperatival, just like judgments of morality and, as I noted in the last paragraph, judgments of what is true. In saying that something is true or moral or beautiful, we effectively demand that everyone agree with us. In fact, I think that in *most* cases, the act of ‘judging’ something has an imperatival element, in that we not only expect but ‘demand’ (implicitly) that others agree with us. Admittedly, Kant disagrees. He distinguishes between judgments of taste (and morality) and others:

⁵ Ronald Dworkin, “Objectivity and Truth: You’d Better Believe It”, *Philosophy and Public Affairs* (1996): 97.

⁶ [See also J. L. Mackie’s classic *Ethics: Inventing Right and Wrong* (1977).]

⁷ Kant’s *Critique of Judgement*, translated by Walter Cerf (Indianapolis, 1963), p. 18.

⁸ Kenneth F. Rogerson, “The Meaning of Universal Validity in Kant’s Aesthetics”, *Journal of Aesthetics and Art Criticism* 40.3 (1982): 301.

...Many a thing may be attractive and pleasurable to him; no one cares about that; but if he declares something to be beautiful, he expects the very same pleasure of others, he judges not solely for himself, but for everyone, and then speaks of beauty as if it were a property of things. Hence, he says, the *thing* is beautiful, and does not count on others agreeing with his judgment of pleasure because they did so occasionally in the past; rather he demands this agreement from them. He censures them if they judge differently and denies them taste, which he yet demands they should have.⁹

His distinction, however, between attractiveness and beauty seems spurious. As far as I can tell, the difference between them isn't of kind but of degree: in saying that an object is attractive I am saying much the same thing as when I say that it's beautiful, the only difference being that I'm a little less committed to the former assertion than the latter, because my 'aesthetic admiration' is more pronounced with respect to the latter. If my expectation/demand that other people (will) agree with me is less pronounced in the case of attractiveness, it is simply because I'm less committed to the aesthetic worth of the given object than I am if I say that the object is beautiful. I can, after all, intelligibly disagree with and argue against someone who denies that a given object is attractive, as if *demanding* that he agree with me; it's simply not very likely that I will, because the aesthetic worth of the object doesn't strike me with as much force as it would if I said that the object is beautiful.

The reason I harp on this is that I want to say that the imperatival element in certain assertions (in *most* assertions) comes partly from the form of declarative sentences themselves, which posit their truth and their own truth-conditions. Even when the speaker doesn't act as if he expects or demands that people agree with his utterance, I want to say that he has the linguistic 'right' to do so—that, given the necessarily truth-positing form of declarative utterances, it's perfectly natural and linguistically justified for a speaker to expect that others will agree with his assertion. For the form of his assertion posits its objectivity—its truth independent of the speaker's attitude—and something that is objectively true is, as I noted a moment ago, thereby supposed by everyone to be a reasonable and virtually obligatory object of belief.

⁹ Quoted in *ibid.*, p. 304.

Before proceeding, I ought to return briefly to the point that Kant makes above about pleasure. He says that if I say something is pleasurable, I don't thereby expect or demand the same pleasure from others. My pleasure exists only for me; it isn't universalizable, as judgments of taste are. –First of all, I think he's wrong. If I judge something pleasurable and somebody disagrees with me, I may well get frustrated with him and say “How can you possibly disagree?! You don't think massages are pleasurable?! What's wrong with you?!” But insofar as Kant's point has plausibility, it is because saying that something is pleasurable is to say that it *gives me* pleasure. It is thus different from a statement like “It is attractive” or “It is beautiful”, for these statements cannot be rephrased as “It gives me attractiveness” and “It gives me beauty”. In short, the subjective, relativistic standpoint is *built into* the assertion that something is pleasurable. And yet even then we often demand that people agree with us that an object is pleasurable! (That is, we often ‘universalize’ our sentiments about what is pleasurable.) This fact is significant.

Part of its significance is that it supports one of my arguments, namely that implicit in the ‘representational’ form of assertions is objectivity and, ultimately, ‘universalization’ (a “lay[ing] claim to everyone's assent”). But there is also a psychological significance, which harmonizes well with the philosophical one. For the example above hints at a broader phenomenon, viz. that a person spends his whole life in essentially the infantile state of mind, according to which “everyone must be like me”.¹⁰ His experience is, naturally, the center of his world, which naturally leads him to interpret it—almost unconsciously—as the center of *the* world, at least the human one. People are, quite literally, self-centered; they therefore rarely understand people with experiences and traits substantially different from their own. And they unconsciously project their own attitudes and beliefs onto (into) other people. So it is predictable that they often call people irrational who disagree with them—for anyone who has the same experiences *they* have (e.g., the same experience of what a massage feels like) should apparently draw the conclusions *they* do (that massages are pleasurable). If they don't, it would seem that they are failing to make logical connections, draw logical inferences.

Thus, just as the nature of language commits us to the objective truth of many of our utterances, so the nature of experience implicitly commits us to the objective, ‘universal’ truth of our beliefs and attitudes. We may not recognize this on a conscious level, but our unconscious or

¹⁰ See “Everybody Must Be Just Like Me”, by Elizabeth Lloyd Mayer, in *Female Sexuality: Contemporary Engagements*, edited by Donna Bassin (Northvale, NJ: Jason Aronson, Inc., 1999), pp. 377–403.

half-conscious commitment to the objective truth and universal validity of our attitudes is revealed in the fact that we tend to get angry at, and treat as irrational, people who disagree with us. We treat it as a kind of imperative that others agree with us—albeit an imperative that is easily forgotten or ignored in social life, since life requires that we get along with people who are different from us. And while in the immediacy of social interaction we may not frown on mutual disagreement—may even encourage it—in the reflectiveness of solitude we may well have disdain for, or be bewildered by, people who don't hold our views on things. We will most likely consider them somewhat irrational.

But what is rationality? Earlier I made a few suggestions about it; but now it's time I addressed the question in more depth. Philosophers usually distinguish between theoretical reason and practical reason: the former is concerned with standards for seeking truth (about the way the world is), the latter with “rational standards that apply directly to conduct”.¹¹ Or, as Carl Wellman says, “practical reasoning is using and following practical arguments, arguments that bear fairly directly on practice, on the pursuit of goals and the choice between alternative actions”.¹² Despite the ambiguity of the words ‘fairly directly’, this formulation may seem clear enough. Theoretical reason relates to what *is the case*, while practical reason relates to what *ought to be done*. Nevertheless, several philosophers have argued that the distinction between the kinds of reason is illusory—or at least, not as important as one might think.¹³ Before continuing, therefore, we should clarify what we're talking about.

Consider the question of whether to classify instrumental reason as practical or theoretical. The instrumental principle states that one ought to take the means to one's ends; instrumental reason relates to the question of what are the best means to achieve a given end. The instrumental principle itself appears to qualify as practical reason, since no amount of word-twisting can reformulate it as an answer to the (theoretical) question of what is the case. Of course, one can say, “It is the case that people ought to take the means to their ends”, thus apparently assimilating the principle to theoretical reason, but this is pure sophistry. Edgley remarks, rightly, that the move fails because “answers of the sort ‘So-and-so ought to be done’

¹¹ Christine Korsgaard, “Skepticism about Practical Reason”, in *Moral Discourse and Practice*, edited by Darwall, Gibbard and Railton (New York: Oxford University Press, 1997), p. 373.

¹² Carl Wellman, “The Justification of Practical Reason”, *Philosophy and Phenomenological Research* 36.4 (1976): 531.

¹³ See, e.g., R. Edgley, “Practical Reason”, *Mind* 74.294 (1965): 174–191.

are not answers to the question ‘What is the case?’...’¹⁴ But because it strikes us as irrational not to take the means to one’s ends, we are inclined to say that the instrumental principle is a norm of reason; and since it isn’t theoretical reason, it must be practical.

On the other hand, instrumental *reason* is surely a manifestation of theoretical reason, in that it answers the question “What is the case?” In order to bring about a certain state of affairs, logic and induction tell us that so-and-so has to be done. Value-judgments do not come into play here; we are asking only what causes will produce what effects. We are logically and inductively inferring conclusions from premises, and doing nothing else; we are asking what *is*.

Prima facie, then, there are two basic categories of reason. So what is the connection between them? Let’s look at the instrumental principle more closely, since it is, so far, the only example I have given of practical reason. What is the source of its normativity? How can the principle that one ought to take the means to one’s ends have normative force? For, as Christine Korsgaard points out in her paper “The Normativity of Instrumental Reason”, the principle construed a certain way is an analytic truth. “To will an end just is to will to cause or realize the end, hence to will to take the means to the end. This is the sense in which the principle is analytic.”¹⁵ In other words, committing yourself to realizing an end logically entails committing yourself to taking the means. But if so, how can it make sense to *prescribe* that people ought to (try to) take the means to their ends? That follows tautologously from their commitment to their end. It’s like saying, “You should do what you are determined to do”, which sounds odd, to say the least. Korsgaard uses some verbal legerdemain to avoid the natural conclusion that the instrumental principle is not really normative, but I think she fails.

An example might highlight the paradox [of how the instrumental principle can seem genuinely prescriptive despite its analytic nature]. Suppose I intend to rent a movie tonight but refuse to go to the video store. “I’m going to buy a movie,” I say to my friend, “but I’m not going to go to the store.” “Huh?” my friend says. He won’t think I’m being sincere, or else he’ll think I’m leaving out a piece of information. No matter how much I insist on my two contradictory intentions, he’ll think I’m just joking. The reason is that it is *psychologically impossible* to commit oneself to an end without committing oneself to the means (the means, that

¹⁴ Ibid., p. 174.

¹⁵ Korsgaard, “The Normativity of Instrumental Reason”, in *Ethics and Practical Reason*, edited by Garrett Cullity and Berys Gaut (New York: Oxford University Press/Clarendon), p. 244.

is, that one knows one has to take). People never exhibit this kind of irrationality, unless, perhaps, they're hopelessly insane.

The solution to the paradox is simple, though. The instrumental principle intuitively seems normative because we ordinarily have a loose understanding of the word 'ends'. We take it to mean something like 'things that one strongly desires'. On this understanding, it's actually a fairly substantive thing to say that people should try to bring about their ends, since what is being said is that they should try to bring about what they *strongly desire*. It's psychologically possible for me to yearn for something while doing nothing to bring it about—but it also seems irrational. Understood in this way, the instrumental principle is both obvious and able-to-be-broken, which is what we intuitively want.

The question now, of course, is why it should be a requirement of reason that one do what one strongly desires. There is, after all, nothing analytic about *this* principle—and since 'abiding by reason' is supposed to mean 'abiding by logic (or induction)', it would appear that by violating this principle I am not violating reason. So perhaps the instrumental principle, understood in the vague way we usually understand it, shouldn't be normative after all. —The way out of these difficulties is to acknowledge that practical irrationality is not a matter of violating only logic. As I said in the beginning, there is a kind of psychological inconsistency which is the true intuitive basis for the judgment that someone is being irrational. If I have explicitly conflicting attitudes (or values, beliefs, intentions, etc.), or I act in ways that imply that I have obviously contradictory values or intentions, I will be called irrational.

Clearly the notion of 'psychological consistency' is vague. Much vaguer than that of logical consistency. Its vagueness, however, is a strength of my account, because it mirrors the vagueness of our ordinary intuitions about when someone is being irrational. Suppose, for example, I'm sitting at home, bored, and I turn on the television. I start watching *Entertainment Tonight*, a show I've always hated. I know there is nothing intellectually or emotionally redeeming about watching it, and I know I'm wasting time, and I know I should be working, and I'm not even really being entertained—yet I keep watching it, inexplicably. The question is, am I being irrational? There are plenty of reasons for me to turn the TV off and apparently none for me to keep it on—but is it right to say that by keeping it on I'm being irrational? One is tempted to say, "In a sense, obviously you are. After all, you're not being motivated by reasons that you acknowledge *should* motivate you. But in another sense, you're being rational. For you clearly

don't want to do work, and there's nothing you'd *rather* be doing than watching TV, and apparently you just don't care that you're wasting time—and perhaps you actually are being a little entertained, despite your refusal to admit that to yourself. So you're being both unreasonable and reasonable.” In this case, intuition hesitates. I'm being both rational and irrational—consistent with myself and inconsistent with myself. Obviously there is a part of me that does prefer to keep watching, as opposed to getting up and doing work; otherwise, of course, I would get up and do work. So in this sense I'm simply doing what I want to do, which means I'm acting rationally. But there is another part of me, as I said, that's disgusted and so forth; inasmuch as I am acting contrary to the desires and reasons of *this* part, I am being self-inconsistent and hence irrational.

The same considerations apply to the example Korsgaard gives in her paper “The Normativity of Instrumental Reason” of the man who, because of a disease he has, must get a series of injections in order to live past 50. He wants to have a long life, but he is also terrified of injections. He's so terrified, in fact, that he refuses to get the treatment, even while knowing that it will mean his early death. “He agrees that a long and happy life is a greater good than avoiding the injections, but he still declines to have them.”¹⁶ Intuition hesitates here in the same way it did in the other case: from one perspective, the man's fear is clearly causing him to be irrational (to act contrary to what he acknowledges is his greater good); from another perspective, he is being perfectly rational, since *right now* his strongest desire is to avoid the injections rather than have a long life, and he is acting in accord with that desire.

Philosophers discussing the status of reason tend to use as thought-experiments examples, like the foregoing, that involve instrumental reason (or else constitutive reasoning). But I think that the normative foundation of the instrumental principle is really the same as that of the ‘principle’ that, say, you shouldn't destroy a painting you think is beautiful, or you should be honest if you think being honest is good: by making a value-judgment, you posit certain norms on your behavior (which aren't necessarily *logically* related to your judgment); if you flout these norms, you're exhibiting irrationality (insofar as you still hold the value in question when you flout the norms). Particular norms are the *psychological* ‘flip-side’ of particular value-judgments. What establishes the connection between the norm and the value-judgment is, I suspect, the universal properties of intentionality and action, and their relation to value-judgments. I'll return

¹⁶ Korsgaard, *op. cit.*, p. 227.

to this question below. My point right now is that philosophers should focus more on value-judgments in general and less on questions about desire and instrumental reason, because desire, after all, is but one kind of value-judgment. (In desiring something—say, financial success—I am valuing the thought of experiencing whatever it is I desire.) And the instrumental principle is but an example of the fact that value-judgments have prescriptive implications. Practical irrationality, then, means violating the prescriptive implications of one’s value-judgment.

I agree with Korsgaard that a good way to characterize irrationality, whether practical or theoretical, is as the “fail[ure] to be motivationally responsive to the rational considerations available”¹⁷. I simply think that her idea of what qualifies as a ‘rational consideration’ is too narrow. A rational consideration isn’t necessarily a matter of adhering to logic and induction;—or rather, in the sphere of *theoretical* rationality, it is. But in the sphere of practical rationality it has more to do with what I have called the ‘psychological implications of a given value-judgment’—implications that can be logical¹⁸ but need not be. What they have in common with *logical* entailments is that, like logic, they are grounded in the principle of non-contradiction. If the injunction to avoid self-contradiction is virtually the guiding law of theoretical reason, so it is with regard to practical reason. The only difference is that what counts as a self-contradiction is somewhat broader in practical than in theoretical reason.

[...]

Now, David Hume looked at all this in a somewhat different way than I have. He seemed to think that there is no such thing as practical reason. (I doubt he even thought that *instrumental* reason is practical reason, since he knew that such reasoning is a disguised form of finding out what is the case—i.e., what cause will bring about a desired effect.) Reason, he thought, cannot in itself motivate anyone to act. Only desires can motivate, and desires are neither rational nor irrational per se. “Reason can teach us how to satisfy our desires or passions, but it cannot tell us whether those desires or passions are themselves ‘rational’, that is, there is no sense in which desires or passions are rational or irrational.”¹⁹ Even outrageous desires cannot be criticized as irrational. To quote Hume:

¹⁷ “Skepticism about Practical Reason”, p. 378.

¹⁸ For instance, sometimes a value-judgment logically entails another one. I’ll give examples later.

¹⁹ Korsgaard, *op. cit.*, p. 374.

'Tis not contrary to reason for me to prefer the destruction of the whole world to the scratching of my finger. 'Tis not contrary to reason for me to chuse my total ruin, to prevent the least uneasiness of an Indian or person wholly unknown to me. 'Tis as little contrary to reason to prefer even my own acknowledg'd lesser good to my greater, and have a more ardent affection for the former than the latter.²⁰

I have already explained in what sense I disagree with these ideas. I think that our acts may well conflict with what we value on a reflective level—say, if we are overwhelmed by fear and therefore act contrary to the (prescriptive implications of the) value-judgments we're reflectively committed to. This, I maintain, should be called a form of irrationality, partly because we intuitively think it is such and partly because it involves blatant self-contradiction: I am acting contrary to what my reflective value-judgments prescribe. If I'm not really committed to them, then perhaps it isn't quite correct to say that by acting contrary to them I'm being irrational; if I am committed to them, though, then I *am* being (practically) irrational. My irrationality is practical because it isn't a violation of logic or induction but rather of the rules of 'psychological entailment' (which decree, for example, that if I value the thought of not watching *Entertainment Tonight*, I ought to stop watching it)—in other words, because it doesn't relate to what is *true*, as theoretical irrationality does.

From another perspective, however, Hume is right that practical irrationality is impossible. For every act (and intention) is the expression of a value-judgment. It presupposes, and can indeed be called the 'actualizing' of, a value-judgment.²¹ For instance, by writing this paper I am 'expressing' or 'demonstrating' my (commitment to the) value-judgment that writing this paper is good, i.e., that I ought to be writing this paper. Similarly, insofar as my watching *Entertainment Tonight* can justifiably be called an *act*, it presupposes the (not fully conscious) value-judgment that watching it is good, at least in the circumstances in which I find myself at that moment. In short, as an agent—i.e., as someone with free will, which means that a given act

²⁰ David Hume, *A Treatise of Human Nature*, L. A. Selby-Bigge, ed. (London: Oxford University Press, 1888), p. 415.

²¹ In a more direct sense, of course, it is the actualizing of an *intention*.

is done because one values it or something it will lead to (or is constitutive of)²²—one values, in some way or another (whether implicitly or explicitly), acting in whatever way one is acting in a given moment. Hence, one is necessarily always abiding by the prescriptive implications of the value-judgment that overrules all others in that moment. And this means that one is never being practically irrational.

I suspect, however, that Hume didn't have such considerations in mind when writing the passage above. Instead, he was articulating the implications of the fact that 'ought' cannot be derived from 'is', that a value-judgment/prescription is not logically inferable from what is the case. He may have been the first person to discover this truth, but he was by no means the last to emphasize its importance. His point anticipates, for example, Kierkegaard's insistence that every decision, every act—even the act of accepting a conclusion validly derived from premises (as well as the act of thus deriving it)—involves an element of freedom, a "leap of faith".²³ Hume doesn't explicitly focus on the element of freedom, but he clearly has it in mind (however obscurely) in denying that 'ought' is derivable from 'is'. He is insisting that there is a gap, a gap of freedom—of non-rationality—between a reason (properly so-called)²⁴ and a value-judgment, which is basically equivalent to saying that there is a gap between a reason and the carrying-out of an act, because an act is the actualizing of a value-judgment (or its prescriptive implications). Therefore, acting contrary to theoretical reason (or to what is true) is impossible. Meaningless. Since Hume thought that theoretical reason is the only form of reason, he thought that it's impossible to act contrary to reason, i.e., truth.²⁵ –To say it one more time, the reason I partially disagree with him about the possibility of practical reason, even though I agree with him that 'ought' cannot be derived from 'is', is that I recognize a form of reason that is not really

²² The reason why the idea of self-control entails *that* idea is that it entails that what I do is done because I want to do it, i.e., because I value it or something it is a means to (or constitutive of).

²³ Kierkegaard thus anticipated Lewis Carroll's argument in his famous paper "What the Tortoise Said to Achilles", in *Mind* 4, No. 14 (April 1895): 278-280. Carroll pointed out the difference between the *acceptance* of a logical rule (such as *modus ponens*) and the *application* of it. The latter is a second act, different as such from the first act of accepting the logical rule in question. One can accept the logical rule but refuse to apply it, because every act, according to Carroll and Kierkegaard, involves non-rationality, or freedom.

²⁴ By 'reason' here I mean a conclusion derived from logic or induction.

²⁵ He recognized, of course, the possibility of theoretical irrationality—but, as will become evident three paragraphs later in the text, he didn't really have the 'right' to acknowledge the existence of theoretical irrationality. For this kind of irrationality is just a matter of the motivational force of an end (namely, knowledge) failing to transmit to the means necessary to achieve the end (namely, adhering to logical and inductive rules). In other words, it is a kind of instrumental irrationality, which itself is a manifestation of *practical* irrationality. So theoretical irrationality is a manifestation of practical irrationality. But since Hume rejected the possibility of the latter, he arguably had no right to accept the possibility of the former.

assimilable to the category of ‘is’: it consists in the fact that insofar as I value something, I psychologically (or perhaps it’s better to say *phenomenologically*) commit myself to acting a certain way. As Korsgaard might say, *I prescribe a law (or laws) for myself*. If I value Schubert’s piece *Erlkönig*, I implicitly ‘prescribe’ that I will not throw away the CD it’s on, that I will not go around telling people it’s a terrible piece, etc. For any given value-judgment (whether implicit or explicit), there are a number of normative implications. If I violate any of them—and if my violation doesn’t grow out of another value-judgment that has temporarily trumped the earlier one—I’m being to that extent irrational. But this kind of irrationality is not a violation of logic or induction; rather, it stems from my acting contrary to the nature of action itself. In violating the normative implications of my value-judgment, I have ‘negated’ (contradicted) the value-judgment itself—even as I still profess to be committed to it—and since the act of valuing is really the basis, and meaning, of action itself (in that implicit in every act is a valuation, and every act is a realizing, a manifesting, of a particular valuation), I have ‘negated’ or contradicted *action*. My two relevant acts (the valuing and the ignoring of the normative implications of this valuing) are mutually vitiating, as well as—in combination—contradictory to the universal meaning of action itself.

In any case, I noted earlier that it seems almost psychologically impossible to act in this way. It is scarcely conceivable that a person could contradict himself in the way I imagined him to in the ‘renting a movie’ example above (as well as in the ‘valuing a painting but destroying it’ example²⁶). At most, *partial* practical irrationality is possible—as when I keep watching *Entertainment Tonight* despite my belief that my ‘greater good’ would be to get some work done.

As I was saying, though, the element of freedom, of non-rationality, in action, which is also that of value-judgments, entails that ‘ought’ cannot be derived from ‘is’. For instance, one cannot legitimately do what Schopenhauer did in justifying an ethics of compassion through the metaphysical consideration that everyone is part of an undifferentiated thing-in-itself behind appearances (and so by treating people compassionately we’re acting in harmony with our mutual oneness). –Or rather, one can ‘justify’ one’s morality however one likes, but (1) if the justification proceeds directly from what is to what ought to be, then there is no logical

²⁶ It is of course possible to value something and yet treat it badly, but in doing so, one shows that he values something more than the thing he’s treating badly. He needn’t necessarily *think* he values it more, but inasmuch as a person does what he wants to do, and what he wants to do is what he values at that moment, *this* person (the one who doesn’t think he really values what he is doing, even though he has chosen to do it) is deluding himself.

connection between it and the morality, and the justification fails; or (2) if it deduces the morality from a more basic set of values, then the logical leap (from ‘is’ to ‘ought’) occurs in the justification of the more basic set, and so, again, the justification of the morality fails. That is, it’s conditional upon the acceptance of the set of values from which it is deduced—and a conditional justification isn’t what one wants in a justification of morality. I’ll return to this point later, but the relevant conclusion to draw from the gap between ‘is’ and ‘ought’ is the one that Hume draws in the passage above: desires, or more generally values, are not criticizable as violating ‘truth’ (whether logical, scientific, metaphysical or whatever).

But given the gap between ‘is’ and ‘ought’, how can theoretically rational considerations motivate? And how is it possible to be theoretically irrational? The answer is that one’s responsiveness to theoretical reason must come from another source, albeit a source that decrees that such reasons will be capable of motivating. For example, suppose I value knowledge. I desire to know truth. My value itself, like all values, is neither reasonable nor unreasonable; it’s just a brute psychological fact. But since I know that the way to pursue knowledge is to carry out logical and inductive processes of reasoning, I will be motivated by such reasons. In themselves they can’t motivate, but given my value, my desire—my ‘end’, namely knowledge—they can motivate. Actually, I’m basically adhering to the instrumental principle by being motivationally responsive to theoretical reasons: careful reasoning is related to the attainment of knowledge as means to end. And as philosophers know, instrumentally rational considerations are able to motivate because the motivational force of the end gets transmitted to the means necessary to achieve the end. –This shows, incidentally, that theoretical and practical reason are very similar, despite my efforts to differentiate them. For the instrumental principle is a manifestation of practical reason. And so, oddly enough, the motivational force of theoretically rational considerations is grounded in practical rationality (in the imperative—which, as Korsgaard says, is constitutive of action as such²⁷—to respect the normative implications of one’s value-judgments), even as the *meaning* of *practical* irrationality (viz., inconsistency with oneself) is but an extension of the meaning of theoretical irrationality. Both kinds of reason presuppose the other.

[...]

²⁷ See, e.g., “The Normativity of Instrumental Reason”, p. 249.

So, how does one go about justifying a particular act? One cannot appeal solely to what is the case, because there is no logical connection between ‘ought’ and ‘is’. Rather, one has to appeal to what is the case *in combination with* other values that one holds. (More precisely: one’s other value(-judgment)s will themselves be included in the consideration of what is the case, since, after all, they exist.) Much of the time this justification will proceed along instrumental lines. I can justify eating at McDonald’s by invoking the facts that I’m hungry, i.e., that I value the thought of eating food right now, that the prescriptive implication of this value is that I find a place to buy food, and that McDonald’s is the only place nearby. –Justifications involving the instrumental principle are easy and unproblematic.

Suppose, on the other hand, you want to justify your approval of a particular movie. How should you do this? First you have to note that you have certain other relevant values, such as a complicated plot-structure, rich characterization, thematic significance, and lush cinematography. Then you argue that the movie manifested these qualities. If you’re right—if you do a good job of arguing that the movie epitomized your values, such that, perhaps, you convince other people to agree with you—then the value-judgment you set out to justify is justified. Notice, though, that it’s only conditionally justified. It wouldn’t necessarily be irrational for someone to reject the values that provided the basis for your approval of the movie. It *would* be irrational if this other person shared with you an even more basic value, from which the values he has rejected can be easily derived.²⁸ This circumstance is unlikely, however. It is more likely that he happens to have different values than you, in which case he isn’t being irrational because he isn’t failing to be motivated by rational considerations. If you don’t have a relevant value in common with him, from which the four values I mentioned above can be deduced, it is senseless for you to try to convince him that he’s wrong or irrational. For values can be deduced only from other values, not from what is the case.

If the two of you do get into an argument about whether the movie was good, it’s likely that you’re really trying to convince him that, despite his denial, he actually does have the relevant values that you do (such as lush cinematography). He just doesn’t realize it at the

²⁸ An example of such a value is music that sounds gloomy. Given that minor modes sound gloomier than major modes, it’s inferrable from the judgment that ‘gloominess is good’ that pieces in a minor key are preferable to pieces that aren’t. This second value, then, can be derived from the first. If someone holds the first but not the second, he is contradicting himself and is therefore, at least *prima facie*, exhibiting irrationality.

moment, so you're trying to coax him out of his momentary confusion or ignorance of his own values.

—I could continue embellishing all these ideas, but the purpose of this paper isn't to clarify rationality. I've been discussing it mainly so as to have firmer ground to stand on in the following. Hence, I won't investigate the relations between the other meanings of 'rational' and 'reasonable', nor will I show to what extent ambiguities in natural language are responsible for confusions about the notion of reason. (It is easily confused with that of intelligence, which is even *more* ambiguous, as well as with intellectual creativity. And 'irrationality' can be confused with a reliance on mistaken premises, or with deficiencies in such non-rational faculties as aesthetic sensibility.) The immediately relevant point is that in criticizing people for exhibiting irrationality, we assume that one of their most fundamental value-judgments is that it's good to have knowledge, it's good to be justified in one's acts and beliefs. Hence, our criticism has nothing to do with "external" reasons²⁹ or non-motivating "reasons of rationality"³⁰ or anything like that. We're simply criticizing someone for not adhering to the prescriptive implications of *his own* fundamental value-judgment. We're criticizing him for not wholly abiding by something like the 'principle of constitutive reasoning' (that you do whatever it is that constitutes your ends—even if they're largely unconscious, as in this case). That's why our criticism has force: he recognizes that, if it's true, he's failing to act in accordance with one of his own values, and is thus being irrational.

The circularity of this situation is a little paradoxical. We criticize someone for not being motivated by reason, which he should be because, as a human, he wants to be justified in his acts; and he (perhaps) heeds our criticism just because, in acting contrary to his value (of being justified in his acts), he is being irrational. That is, the reason he doesn't want to be irrational is that it's irrational to be irrational. There is no other reason to adhere to reason than the fact that *not* to adhere to it would be contrary to reason. This fact somehow strikes us as bad enough in itself; no other justification is needed for our acting rationally. We value it for its own sake. This is surely related to the fact that rules of reason are *logically* self-justifying too. As James Dreier says, "there is no sense at all to be made of the question of whether we have any reason to follow

²⁹ See Bernard Williams, "Internal and External Reasons", in *Moral Discourse and Practice*.

³⁰ See James Dreier, "Humean Doubts about Categorical Imperatives", in *Varieties of Practical Reasoning*, edited by Elijah Millgram (Cambridge: MIT Press).

the rules of rationality”.³¹ These rules (norms) are therefore unlike other rules, such as rules of prudence or morality, in that they are their own justification. One cannot ask for a reason to act rationally without presupposing norms of rationality—just as one, almost of psychological necessity, is committed to the *value* of being rational(ly justified in his acts).

These ideas have implications for morality, but I’ll postpone that discussion until I’ve finished the analysis of values in general that was interrupted by the analysis of reason. I was in the middle of pointing out that we implicitly ‘assume’, in the act of utterance, that the assertion we make is true. We may recant later, but when we utter it we obviously think it’s true and justified. Language forces this fact on us, as does our own unconscious view that “everyone must be like us” (that our own experience is somehow universal), as does our foundational commitment to rationality. Kant’s focusing on beauty, therefore—on our how we universalize judgments of taste—was somewhat misleading: we posit the objectivity of *all* our value-judgments, and we reason (or argue) about them as if there is an objective truth to the matter. Our basic desire to be right, or at least justified, in our acts is the reason for this. The question is, is it right? Are values and their corresponding norms somehow objective? Can they be ‘true’?

I should describe the realist’s position more fully before answering the question. I’ll quote Korsgaard’s definition: “realism is the view that propositions employing [normative] concepts may have truth values because such concepts describe or refer to normative entities or facts that exist independently of those concepts themselves. We have the concepts in order to describe or refer to those facts.”³² I have said that realism is implicit in our ordinary use of language, because propositions, in positing truth-conditions, represent an independent reality. Realism, therefore, is the commonsense view. It is also the view that people find most attractive, because it (potentially) justifies their certainty about their own values. That is, since people like to think they’re right, a position that holds out the possibility that their value-judgments are objectively true is seductive. Is it correct, though?

It should be evident from the foregoing that I answer that question negatively, except with respect to norms of theoretical reason (and mathematics, which is just a special kind of pure reason).³³ But that’s because these latter norms, unlike others, are concerned not with what is

³¹ Ibid., p. 29.

³² Korsgaard, “Realism and Constructivism in Twentieth-Century Moral Philosophy”, *Philosophy in America at the Turn of the Century* (2003): 100.

³³ The *value* of reason is not somehow ‘objectively true’, but *rules* of reason are.

good but with what is *true*. This unique property is manifested also in the fact that they're categorical: they are their own justification. We can't question them without presupposing them. This is why Aristotle said that in order to win an argument with someone who doubts the principle of non-contradiction, all you have to do is get him to assert something.³⁴ Basic rules of reason, then, are self-justifying, which distinguishes them from all other sets of rules—whether they be 'derivations' from the instrumental principle (like prudential maxims) or norms that are the direct flip-side of value-judgments (like "One ought not to murder", which is the 'flip-side' of the judgment "Murder is wrong"). One can coherently argue for or against these other sets of rules, as one *cannot* argue for basic rules of reason. But prescriptions of instrumental reason,³⁵ while derivative of basic norms of rationality, have more in common with the latter than they do with norms that are the flip-side of value-judgments, for they are concerned solely with how to cause a given effect, while the others are normative entailments of value-judgments.

It is these norms, and the values they presuppose, that cannot be 'objective' or 'true' in the way that norms of rationality are, precisely because they are not self-justifying (or—in the case of 'prescriptions of instrumental reason'—deducible from self-justifying rules in combination with given states of affairs). In other words, it's because there is a gap of freedom between what is and action (i.e., what ought to be) that these norms/values cannot be a part of reality, for they are nothing but actions in the form of 'potentiality'. (Remember I said that implicit in every act is a corresponding value-judgment, and every act is an expression of a value-judgment. Similarly, every value-judgment/norm is an unactualized act. "It's wrong to murder" is an unactualized refraining-from-murder; "one ought to be honest" is an unactualized being-honest. It is a positing-the-act-of-being-honest.) They are, in fact, posited modifications of reality: after all, they are merely what *ought to be*, as opposed to what *is*. By definition, then, values and norms (except norms of reason) cannot be objectively true. Like actions, they are expressions of freedom; and freedom, as Sartre argued, is essentially a negating of the real. In acting (i.e., being free), I am modifying reality, I am asserting myself. I am changing what is the case at that moment. Since value-judgments, and their normative implications, can be thought of as unactualized acts, and acts, in being expressions of freedom, are negations of what is real,

³⁴ See "The Normativity of Instrumental Reason", p. 248.

³⁵ For example: if your goal is to eat food, go to the kitchen and get some food.

value-judgments are negations of objective reality. By definition, then, value realism and normative realism (excluding realism about norms of rationality) are false.

There are many more arguments one can bring against them. Consider John Mackie's famous "argument from queerness":

If there were objective values, then they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe. Correspondingly, if we were aware of them, it would have to be by some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else.[On this view,] an objective good would be sought by anyone who was acquainted with it, not because of any contingent fact that this person, or every person, is so constituted that he desires this end, but just because the end has to-be-pursuedness somehow built into it...³⁶

Values motivate the person for whom they are values. If they were objective, as $2 + 3 = 5$ is objective, or as "In order to eat food, one must put it into one's mouth" is objective, they would still have to motivate. But neither of the propositions I just mentioned *motivates* (the second might seem to, but it actually derives its motivational force from the *value* attached to the thought of eating)—and, in general, I can't conceive what it would be like for an objective truth to motivate. The idea doesn't make sense. But of course this fact is predictable, given that value-judgments and acts are *negations* of the real, i.e., assertions of freedom. For something real (true) to have to-be-pursuedness built into it would be a negation of its realness, since action is a negation of what is. So Mackie's argument from queerness is actually an argument from self-contradictoriness.

Or suppose we talked to a classical utilitarian, who would argue that "an agent is obligated to perform an action when there is a rule specifying that actions of that kind are to be performed. [This rule, according to the utilitarian, is] the principle of utility. It is because of the existence of this rule that we characterize actions as obligatory or forbidden." The problem is that one could intelligibly say, "Why should I care if there's such a rule? What's it to me? It's just some external norm that I reject." Or, as Korsgaard says, "Are we obligated to obey the rule?"

³⁶ From *Ethics: Inventing Right and Wrong*, in *Moral Discourse and Practice*, pp. 96, 97.

If one is obligated to obey the rule, then the notion of obligation exists prior to the existence of the rule. We cannot explain obligation in terms of the rule, as something that arises from it. On the other hand, if we are *not* obligated to obey the rule, then it seems we may permissibly ignore it, and so we have not after all explained why the actions it directs are obligatory.”³⁷ This argument works against a great many moral theories. I’ll explain why in a moment.

Another consideration against realism is that human beings currently have no obvious criteria for what would constitute a truth about values or norms. Science and philosophy have criteria for determining truth, but what sort of criteria could we use to decide whether a moral claim is true? Perhaps agreement among everyone? But this fact in itself would have no bearing on moral truth. The implicit premise that the more people agree with a claim, the greater is its likelihood of being true is not necessarily plausible. Moreover, it is perfectly conceivable that nearly everyone could agree that a certain state of affairs is morally acceptable even though it isn’t. Think of slavery in ancient Greece, or in the American Confederacy. In any case, it’s a vast logical leap to conclude that murder is wrong from the fact that all cultures *think* it’s wrong. But what other criteria are there? –The point is that, if there isn’t a way to confirm its truth, it is unclear what is being said in saying that a value-judgment is true.

Notice that no *argument* can establish the truth of a given value-judgment. Suppose I say, “Raphael’s painting ‘The School of Athens’ is beautiful”, and someone asks me why I think so. I might say, “It has a certain symmetry, its use of perspective is realistic, and it shows an incredible amount of detail”. But my interlocutor can simply deny that these factors entail its beauty. He can say, with G. E. Moore, that it’s an open question whether symmetry etc. is beautiful. If he were philosophically sophisticated he might invoke Kant: judgments of taste are not logical per se, i.e., not based on conceptual relations, but intuitive and subjective. One has to *experience* the object in order to decide whether it’s beautiful; one can’t judge (or defend) its beauty simply by considering its properties in abstraction from the object—i.e., by considering a *list* of them.³⁸ However, the ‘open question’ argument applies to *all* value-judgments, not only judgments of taste. If someone says that an act was good because it demonstrated sensitivity for

³⁷ Korsgaard, op. cit., p. 111.

³⁸ See Frank Sibley’s excellent article “Aesthetic Concepts”, in *The Philosophy of Art: Readings Ancient and Modern*, edited by Alex Neill and Aaron Ridley (New York: McGraw-Hill, 1995), pp. 312–331. He argues that aesthetic properties, even non-evaluative ones like ‘graceful’, ‘delicate’, ‘elegant’, and ‘garish’, are not condition-governed or rule-governed. They’re intuitive: we apply them to an object on the basis of how it intuitively strikes us. –Kant was onto something when he contrasted the aesthetic faculty with the cognitive.

a person's feelings, it's intelligible to ask whether such sensitivity is itself good. This fact is simply another manifestation of the gap between 'is' and 'ought'. One cannot deduce from an object's (or action's) 'natural' properties that it is *good*.

On the other hand, if I justify my value-judgment by invoking other evaluative words— if, for example, I say that Raphael's painting is magnificent, sublime, awe-inspiring, and hence beautiful—someone can challenge the value-judgments from which I deduced the one I was arguing for. He can say that *those* judgments are problematic, inasmuch as they are made on the basis of the object's possessing certain natural properties (whether they be colors or shapes or whatever). So, once again, the gap between 'is' and 'ought' gets in the way.

To say it differently: evaluative words like 'good' and 'beautiful' and 'unjust' are action-guiding. Which is just to say that a judgment in which they occur is an 'unactualized act'—or acts—of some sort. (I say 'acts' because a value-judgment can have more than one prescriptive, or normative, implication.) Non-evaluative words are not action-guiding. So the difference between the two categories is basically the difference between a factual proposition and an act. Clearly *these* categories are radically different; the one is (supposed to be) a replication of reality—a reproduction of it through the medium of language—while the other is a modification of reality, an intruding-of-oneself into reality. To say, then, that one is 'implied' by the other—that an act, or a value-judgment, is implied by a truth—is virtually meaningless. What would it be like for an act to be logically entailed by something (aside from another act or a value-judgment)? Acts are essentially free. They can be *causally* necessitated—say, by the firing of particular neurons—but, first of all, this isn't logical necessity, and second of all, acts are so necessitated only insofar as one abstracts from their *act-ive* character.

But if acts cannot be 'logically entailed' by a truth, they cannot be truths themselves. (A truth, after all, entails itself.) This accords with ordinary usage of the word 'truth'. We don't say that an act is true, only that a proposition is. But if acts are value-judgments 'actualized', then such judgments themselves cannot be true. They aren't the sort of things to which the property of 'truth' applies. It certainly *seems* they are, though. The form in which they are asserted posits their truth. But if this is so, their form contradicts their content. Their content cannot be true—just as the content of an act cannot be true—but their form requires that they be capable of bearing truth-values. This conflict between form and content means that, in a way, they are self-contradictions. Not in the way that "The cat is not a cat" is a self-contradiction, however: *that*

statement has a self-contradictory *content*. The content of a value-judgment contradicts not itself but the form in which it is expressed—the ‘representational’ form, the form of representing-an-independent-reality, which is implicit in declarative utterances.

What this self-contradiction means is that, just as “The cat is not a cat” is in some sense meaningless, so are “Lying is wrong”, “The painting is beautiful”, and “One ought not to murder”. These statements don’t seem as meaningless as “The cat is not a cat” because their content contradicts not itself but its form, and so they aren’t as blatantly self-contradictory as “The cat is not a cat”. Still, they have an element of self-contradiction; and it is partly this fact that accounts for the intuition that, when someone says, e.g., “Lying is bad”, or “Hitler was evil”, it isn’t entirely clear what he is saying. The meanings of ‘bad’ and ‘evil’ are....strangely meaningless. For these words too, like the value-judgments in which they’re embedded, implicitly ‘represent’ something else, in this case a concept that us supposed to have some kind of descriptive content;—in other words, to a property of certain things in the world. Hitler had the property of being evil, lying has the property of being bad. Properties are supposed to exist in the world, just like the truth-conditions of propositions. (For example, colors are ‘out there’ in the world, as are shapes. We even implicitly interpret *sensations* as being ‘in’ the objects that cause them: an ice cube *is* cold, fire *has the property* of being hot.) But an object or act cannot *have the property* of badness or goodness, because values are merely the essential components in value-judgments, which are acts in the form of ‘potentiality’—and an act isn’t the sort of thing that can be a ‘part’ of something in the world. It’s inherently a *changing* of the world. Therefore: values, like value-judgments, contradict themselves. Their content contradicts their linguistic form.

But what exactly is the content of a value or a value-judgment? I have said it is an unactualized act, but this is vague. A better way to say it is that a value-judgment is essentially the speaker’s subjective reaction to the given object projected onto the object—imaginatively made into a property of the object. But the projected thing isn’t precisely the reaction itself; rather, it is the property of ‘causing this particular reaction’, the reaction that the value-judgment manifests. The speaker effectively reifies his reaction and projects its ‘objective referent’ (so to speak) into whatever it is that he is judging. If he says that lying is bad, he is projecting into the act of lying some such property as ‘inherently causing disapproval’. This, however, is really just the property of ‘inherently not to be done’, which is why my two accounts of the content of

value-judgments amount to the same thing. For ‘inherently not to be done’ is a posited action, an unactualized action (or, in this case, *refraining* from action—which is itself, of course, action).

I suppose this second account is a version of projectivism. We project the ‘causal flip-side’ (or ‘objective referent’) of our reaction onto objects, treating it as if it adheres in the object like any other property. The wrongness of murder is (supposed to be) what explains my belief that it’s wrong; the beauty of a painting is what causes [or justifies] my belief that it’s beautiful. In fact, of course, there are features of the object, in combination with my psychological and social conditioning, as well as my biology, that give rise to my value-judgment; however, these features are not what I unthinkingly think they are when I utter my judgment. They aren’t wrongness or beauty or ugliness or goodness or whatever; they’re facts that science can explain, such as the effects that certain visual or auditory stimuli have on my nervous system, etc. Since science and psychology can explain the causal mechanisms through which I make value-judgments, there is no explanatory role for any kind of realism in the sphere of norms and values (excepting norms of rationality). All vestiges of realism are merely redundant and implausible relics of “common sense” and ordinary language use.

What exactly the “reactions” I’ve mentioned consist of has been a matter of debate among philosophers. Are they beliefs, feelings, preferences, attitudes, or inclinations? (I think that really they can be all five.) Whatever they are, though, the value-judgment/normative statement is in some sense an expression of them. The expressivists, then, are partly right—as are, incidentally, most other schools of ethics. Noncognitivists, cognitivists, realists (but only insofar as they describe the attitude implicit in our ordinary value-judgments), emotivists, intuitionists, “sensibility” theorists (such as John McDowell³⁹), and constructivists. Construed in a certain [limited] way, none of these positions precludes the others. They merely emphasize different aspects of what it means to make value-judgments. After all, no type of act is a simple affair, something that can be adequately explained from only one perspective. There are multifarious intentions involved, behavioral dispositions, intuitions, logical implications and presuppositions, physical movements, perceptions of objects and reactions to them, neural mechanisms. Even the phenomenology of a single act is an extremely complex thing. Especially when you recall what I argued in the first part of this paper—that whether a sentence expresses a value-judgment depends on the context in which it’s spoken, and that utterances can have

³⁹ See his “Values and Secondary Qualities”, in *Moral Discourse and Practice*.

evaluative features without really expressing value-judgments per se, and that sometimes it's impossible to tell whether a given utterance expresses a value-judgment or a purely *descriptive* proposition—you'll realize that the truth cannot be *only* projectivism or *only* constructivism. Maybe one theory does a better job of explaining *this* particular utterance than *that* one, such that a different theory is more appropriate for the latter (because its context differs from that of the former). Open-mindedness should reign in every intellectual field, but especially in philosophy.

I won't show the precise ways in which all these theories can be reconciled, since this paper is supposed to be only a sketch and it's already too long. Instead, I'll argue in more detail about why morality in particular cannot be meaningful, or at least as meaningful as it's taken to be.

Everyone agrees that something about morality distinguishes it from other spheres of values and norms. Most importantly, it seems as if moral values are somehow more 'absolute' and 'obligatory' than, say, aesthetic values. This intuition is manifested in the conviction that there is such a thing as moral *rightness* and *wrongness*, as opposed to mere *goodness* and *badness*. Ugliness, for instance, is bad, but few people would call it *wrong*. That word sounds too forceful for aesthetics, too categorical. On the other hand, murder is indeed called 'bad', but more often it's called 'wrong'. So moral values are supposed to be more binding than aesthetic values. In fact, their obligatoriness differs in kind, not in degree.

Another way to say this is that moral norms are *duties*. They're *imperative*. This criterion is what we use to distinguish mere social norms of how to treat others from moral norms. For example, kindness isn't usually considered a moral value, because behaving kindly is merely *good* rather than *right*. It isn't a duty. Not to lie, however, is considered a duty, so it's a moral norm. Insofar, then, as a given type of social behavior doesn't seem intuitively to be *imperative*, or *right*, it doesn't fall under the heading of 'morality'. It's merely a common run-of-the-mill value.

So morality, by definition, consists of duties. This doesn't, by the way, foreclose any possibilities about the *content* of morality. It doesn't, for instance, rule out Carol Gilligan's "ethics of care", which is supposed to be contrasted with Kant's deontological system. Care, compassion, empathy: nothing in principle rules them out as maxims of morality. They simply have to be called duties. It has to be a *duty* that people act with compassion and empathy.

Otherwise such acts can't properly be called moral. –This conception of morality isn't a stipulation; it's implicit in the way we talk.

The question is, what is a duty? The most obvious way to define it is in terms of other normative words. A duty is something that people are obligated to follow; it's an imperative. But what does *that* mean? Evidently if we define that in normative terms, we'll run into the same problem again. Its meaning won't be clear. –This problem, of course, is just an implication of what I've been arguing. Normative and evaluative terms are not fully meaningful, given their (semi-)self-contradictory character. They implicitly pick out some independent reality even as their only real content is their action-guidingness, their positing of an act—which amounts to their reification of an attitude of approval or disapproval (in the person who speaks them). They implicitly project their action-guidingness into the world, which is a contradiction because action is inherently free, i.e., reality-negating. For a posited act (i.e., 'action-guidingness') to be *in* the world is a total contradiction, virtually meaningless. Therefore, normative and evaluative words intuitively strike us as somewhat lacking in content, though we have trouble pinpointing exactly *how* or *why*. But I have clarified the matter, at least in outline.

Again, these facts explain why any definition of 'duty' in terms of *non*-normative words must fail.⁴⁰ Such words, by definition, lack the element of action-guidingness, which means they fail to capture the force of normative words. G. E. Moore's 'open question' argument, I have said, consisted of nothing but a muddled intuition of this fact—i.e., of the difference between 'is' and 'ought'. 'Ought' is self-contradictory (except when it's an 'ought' of *reason*), while 'is' is not; hence, the former cannot be defined except in terms of other normative/evaluative words, which capture both its element of self-contradiction and its action-guidingness.

There is, however, one substantive way to define 'ought': namely, the way in which I defined the 'ought' of reason. Reason justifies itself—its basic norms are self-justifying—and any imperative of reason (such as an imperative that grows out of the instrumental principle) is ultimately justified by, or deduced from, rules that are self-justifying [or self-evident]. To say, therefore, that one ought to act in a certain way because reason dictates it is just to say that the

⁴⁰ E.g., "any act that adheres to the principle of utility is a duty". There are of course other, unrelated, problems with this utilitarian definition, but the fact that it falls victim to Moore's argument—which, by the way, basically consists of asking the question "Why?" ("Pleasure is the good." "Why?" or "The good is what we desire to desire." "Why?", etc.)—is itself fatal to it. In any case, this definition doesn't explain what the word 'duty' *means*; it says only which acts are duties. But the task that has to be preliminary to enumerating the various duties is clarification of 'duty' itself.

act is deducible from fundamental principles of logic in combination with a description of the relevant state of affairs; it is, in other words, to say that *this is the way reality is*. To achieve an end, one *must* act in such-and-such a way; to do this mathematical equation correctly, one *must* add so-and-so together. ‘Oughts’ of reason are actually ‘oughts’ of ‘is’, which is why they are not meaningless. They are deducible from self-justifying rules.

It was Kant’s genius to understand all this without really understanding it (or, at least, its anti-moralistic implications). This is why he defined duties as principles of reason, as self-justifying maxims of practical reason. One ought to follow a duty just because it’s a duty—that is, because it’s self-justifying. His answer, then, to what a moral imperative is and why we ought to follow it is that it’s a *categorical* imperative, like imperatives of reason, and hence its truth is self-evident. Like logic, it doesn’t depend on our subjective desires or our historical circumstances. It’s an *a priori* truth, like $2 + 3 = 5$.

Kant’s answer to what moral imperatives mean is the only possible substantive one,⁴¹ the only one that neither defines ‘duty’ circularly (in terms of other problematic normative words) nor falls victim to the ‘open question’ argument. Indeed, any answer that succumbs to Moore’s argument (i.e., the is/ought dichotomy) is flawed not only for that reason but also because it fails to answer the main question, which is how the concept of ‘duty’ (or ‘moral imperative’) *itself* should be defined. Theories that succumb to the is/ought distinction consist in saying that such-and-such an act is a duty. Hobbes’s social contract theory is an example, as is Bentham’s utilitarianism. Rather than giving an intuitive description of the meaning of ‘duty’, they stipulate that certain acts are duties. In other words, they define the concept extensionally rather than intensionally. The intensional definition, which is what we intuitively want, has to be something like ‘a self-justifying maxim of action’, ‘an intrinsically right rule to live by’. It has to make some reference to reality or objective truth, because, given the representational nature of language, our intuitive understanding of normative words involves some reference to ‘truth’. Realists like Samuel Clarke and Richard Price⁴² understood this, but Kant gave the definition in question its most sophisticated form.

⁴¹ Excepting the theological one, which defines duty as anything that God commands us to do. To abide by a duty, then, is simply to do God’s will. This would be a good definition if it weren’t for God’s nonexistence. (Also, the idea of God is full of contradictions.)

⁴² See, e.g., Martha K. Zebrowski, “Richard Price: British Platonist of the Eighteenth Century”, *Journal of the History of Ideas* 55.1 (1994): 17–35.

It isn't necessary to discuss Kant's answer to what the proper formulations of the categorical imperative are, because there are fatal flaws even in his framework. The first arises from the fact that norms of reason, while categorical in themselves, are hypothetical with regard to our explicit *application* of them. A person need not and does not follow them at all times—indeed, this is what makes (theoretical) irrationality possible—because, as I argued earlier, rules of reason are not intrinsically motivating. They can't be: action-guidingness cannot be a part of reality or truth(s of reason). It comes from values. Therefore, we guide ourselves according to (conclusions derived through theoretical) reason only insofar as we value truth.⁴³ The value transmits its motivational force to the rules of reason, such that we are “capable of being motivated by rational considerations”. But these considerations *in themselves* are not what motivate us;—or rather, they motivate us only on the basis of our valuing them (i.e., their truth), not somehow ‘intrinsically’.⁴⁴ If we don't value them, we don't treat them as categorical. In a sense, then, they're optional, as it's optional to ‘take account of’ *any* aspect of reality.⁴⁵ We do so only insofar as we value that aspect, whether as an end, a means to an end, or constitutive of an end.

Thus, even if the categorical imperative were categorical in the sense that the rules of logic are, our use of it would necessarily be hypothetical. Our intrinsic phenomenological freedom is the reason for this. It's meaningless to say that we *have* to (choose to) follow the categorical imperative, just as it's meaningless to say that we *have* to (choose to) follow the norms of logic. Whether the norms of logic are somehow innate in our brain [i.e. in our cognitive wiring] is another matter; and it may, in this respect, be true that, descriptively speaking, we ‘have to’ follow the rules of logic, just because they structure how we think. But this isn't true of

⁴³ For instance, it can be argued that being Christian is theoretically irrational (though not necessarily practically so), as is being anti-Semitic, as is being prejudiced or superstitious in any way.

⁴⁴ In §344 of *The Gay Science*, Nietzsche emphasizes that even the “scientist”—i.e., the person who prides himself on his objectivity, his commitment to truth—has a fundamental non-rational faith, viz. his faith in the value of truth. This is what motivates him, this faith, or value-judgment, which is, as such, comparable to the priest's faith in God. The section is interesting as revelatory of Nietzsche's quasi-Humean conception of reason. I think that the speculations on the origin of the commitment to truth (he considers it a manifestation of the commitment to morality) are problematic—at any rate, very imprecise—but the section is thought-provoking.

⁴⁵ Dostoyevsky has his protagonist in *Notes from the Underground* rant against the necessity and so-called value of ‘reason’. Freedom is a higher truth than something as ‘necessary’ as mathematics. “...Two times two is four is a most obnoxious thing. Two times two is four—why, in my opinion, it's sheer impudence, sirs. Two times two is four has a cocky look; it stands across your path, arms akimbo, and spits. I agree that two times two is four is an excellent thing; but if we're going to start praising everything, then two times two is five is sometimes also a most charming little thing.” (From the translation by Pevear and Volokhonsky.) Of course, even if we try to reject reason we can't do it completely. Our thought-processes are saturated with the rules of logic—which is also to say that implicitly we ‘value’ them, even if we try not to.

the categorical imperative. Its premise—which is morality’s premise—is the existence of free will. Therefore, our ability to choose whether (i.e., to what extent) to follow reason is also our ability to choose whether to follow the categorical imperative.

Another way to say this is that, while the *truth* of the categorical imperative might be self-justifying, or rather self-evident (which it isn’t, by the way)—as is the truth of rules of reason—the necessity of the imperative’s application is not.⁴⁶ It cannot be, given what ‘application’ means in this context. For it presupposes choice, that is, non-necessity. Thus, to say that the categorical imperative—or, in general, a moral duty—is categorical in the sense required for morality, i.e., categorically applicable, is self-contradictory and basically meaningless. *Nothing* is necessary in this way. Nothing *can* be, because the idea doesn’t make sense. The notion of duty amounts to a conceptual fusion of free will with necessity, which is senseless because the two concepts are negations of each other. –Basically, ‘duty’ is like any other normative/evaluative word, but to an extreme. It involves a reification and projection of *absolute* approval, as contrasted with a form of approval that’s less imperative.

A third way to express the impossibility of moral duties is to return to the point I made above: Kant’s answer to what duty means is the only possible meaningful one. He intuitively understood that the only meaningful form of ‘ought’ is the ought of reason, so he modeled his conception of moral duty on the latter notion. But with regard to reason there is no gap between ‘ought’ and ‘is’, as there is with regard to values. A value is defined by its action-guidingness, which amounts to its negation of what *is*. Since reason is a component in what is—on account of its objective truth—values involve a negation of reason. Insofar as one values something, one is not following pure reason. Since moral duties are also moral values, by their very nature they cannot be imperatives of reason. In short, the only possible substantive answer to what duties are has, ironically, almost nothing in common with what a duty would have to be. Which is why it is, in the end, a virtually meaningless answer.

The failure of Kant’s project is exhibited by the fact that he had to argue extensively for his definition of moral duty. If his definition had really been a self-evident *a priori* truth, like norms of reason, argumentation would have been basically superfluous. (Also, it would have been a mystery why it had taken mankind until the 18th century to articulate the truth.) Moreover,

⁴⁶ Of course, its truth would consist precisely in the universal necessity of its application, but this is exactly why it’s meaningless.

as I've said, his failure was preordained as soon as he had set himself the goal of defending a set of values. No argument that he or anyone else, such as the neo-Kantian Korsgaard, could have put forward would have established his values, or any values, because at some point in any such argument there occurs a leap from what is to what ought to be. Even supposing (what is false) that someone could prove that a particular value logically followed from a truth, we would still have to value truth in order to value that value. So, no matter what, it's always 'permissible' to reject a given value. Nothing can obligate us to follow it, because 'obligation' is meaningless.

One should keep in mind that values and norms (except rational norms) don't have to be genuinely meaningful in order to do their job, which is to facilitate a particular kind of communication. Simon Blackburn has argued that objective-looking normative (or evaluative) predicates, like 'good' and 'ugly', are a necessary prerequisite for "serious, reflective evaluative practice, able to express concern for improvements, clashes, implications, and coherence of attitudes". Without evolving such predicates we would have remained on the level of interjections like "Boo!" and "Yay!", which would have severely hindered complex processes of economic production. Perhaps such production would have been impossible, thus making impossible social evolution. We also could not have articulated 'obligations' between people: rituals like making a promise would have been mute, and thus impossible.⁴⁷ To develop them we had to "invent a predicate answering to the attitude [in question], and treat commitments as if they were judgments, and then use all the natural devices for debating truth".⁴⁸ Only comparatively late in human evolution could we start asking misguided questions like "Is lying *truly* bad?"

But where does all this leave us? If value-judgments are ultimately not rationally justified [or justified purely by reason], and if we know this, how can we go on believing them? After all, in believing them we are believing in their truth. Especially with respect to morality, the truth is far more relativistic than we want it to be. Not only is murder not always bad, but, strictly speaking, it isn't *ever* 'bad'. 'Bad' is meaningless. As are 'evil', 'wrong', 'ugly', 'beautiful', and all other evaluative words. (Words that have a primarily descriptive content but evaluative overtones, like 'interesting', are basically meaningless insofar as they have evaluative overtones. They may also, of course, lack a clear *descriptive* meaning—'intelligent' is an example—but this

⁴⁷ See Nietzsche's *Genealogy of Morals* for speculations on the historical and 'moral' importance of the making-promises ritual.

⁴⁸ Simon Blackburn, *Spreading the Word* (Oxford: Oxford University Press/Clarendon, 1984), p. 195.

is a separate issue.) They do have *some* content, but to the extent that they're supposed to be objective, they're meaningless. Knowing this, how can we with a good intellectual conscience continue making value-judgments? Are we condemned to self-delusion?

To a large extent, yes. Most people will always believe in the objective justifiability of their value-judgments, thus deluding themselves. Those of us who have intellectual integrity, however, are condemned neither to self-delusion nor to cognitive dissonance: we can navigate these shoals by relying on the fact that certain value-judgments are deducible from others. If we value, say, alleviation of suffering more than life for its own sake, then we're being irrational if we aren't pro-choice with respect to the question of abortion (because abortions are done for the sake of alleviating some kind of suffering). Indeed, one can even argue against pro-lifers if one shares a more basic value with them, such as well-being, from which a pro-choice position can apparently be deduced. Since self-conscious inconsistency is irrational, and people value rationality, by convincing them of their inconsistency one may be able to convince them to change their minds about an issue. Most people do, in fact, share basic values, even apart from rationality; and it is rational to argue with them on the basis of these shared values. One's value-system cannot, strictly speaking, be *correct*, but it may be more consistent and to that degree more rational than one's opponent's.

Suppose I want to argue that Tchaikovsky's fifth symphony is better than Beethoven's second. In the light of this paper, the real content of my claim has to be that the former exemplifies the set of criteria I use to judge musical worth better than the latter. I may not be exactly clear on what my criteria are—perhaps I'm just intuitively more moved by the former than the latter, or its harmonies seem more sophisticated, or maybe I have critically analyzed the scores and decided that Tchaikovsky makes more creative use of the sonata form than Beethoven does—but, whatever they are, if you share them while disagreeing with me then we can have a rational argument. If we don't share all of them then I might try to deduce the criteria we don't share from the ones we do, or from other value-judgments we have in common. Ultimately, though, my judgment is nothing but an expression of subjective approval, and as such doesn't have much force. It has force only to the extent that others agree with my criteria. If, say, Beethoven and Tchaikovsky had substantially the same goals in mind when writing their music—e.g., critical approbation, creative manipulation of the sonata form, and avoidance of

bombast or any kind of “florid and superficial” beauty⁴⁹—and if their goals are my criteria, then my judgment bears on the question of whether one composer achieved his purpose more fully than the other. These conditions, incidentally, are not unrealistic. Artists, even between genres, tend to measure their work by the same standards; for example, an aversion to sentimentalism is almost universal—across cultures, too. If the common purpose of art is something like connection with as wide and intelligent an audience as possible, then standards can be rationally justified on this basis (i.e., as means to an end). Sentimentalism is bad because it hinders the spectator’s identifying with the artist and his work; creative fecundity is good because, among other things, it precludes boredom in the spectator.

Indeed, the presence of common purposes and standards makes possible cross-cultural artistic comparisons. Classical Indian music is inferior to classical Western insofar as it is less creative, less instrumentally rich, less ambitious. It isn’t ‘*objectively*’ inferior, though, unless that term means that Western music exemplifies the common criteria more satisfactorily than Indian music. But nothing necessitates that these be the criteria used; one can adopt any criteria one wants. One’s value-judgment, however, will have interpersonal force only to the degree that others share one’s standards (because, again, people value rationality; if they think their judgments are mutually inconsistent, they’ll change them accordingly). —It certainly seems, by the way, that the human psyche has common aesthetic standards virtually hardwired into it, seeing as particular artworks resonate through the ages.

I agree with Ronald Dworkin that ethical and aesthetic relativism is pernicious. “It is now strenuously argued,” says Dworkin, “that since there is no objective truth about interpretation or art or morality there can be no standard of merit or success in artistic or moral or legal thought beyond the interest a theory arouses and the academic domain it secures.”⁵⁰ Postmodernism is a mendacious creed—a philosophical nihilism that criticizes everything but itself. That there is no objective truth about which values are best does not mean mankind can have no universal ethical and aesthetic standards. Given that the vast majority of people have common intuitions about what is valuable, we might as well take these intuitions at face-value and adhere to them. Murder is ‘bad’, theft is ‘bad’; excessive sentimentalism is bad, inauthenticity of any sort is an artistic crime. Prejudice is bad not only morally but also with respect to reason: it entails an insensitivity

⁴⁹ Hume, “Of the Standard of Taste”, in *The Philosophy of Art*, p. 262.

⁵⁰ Dworkin, op. cit., p. 89.

to rational considerations. People across cultures share the intuition that whatever exalts humanity is for that reason good, while whatever degrades or demoralizes it is bad; and they agree, to a large extent, about what it is that thus exalts or debases. “The general principles of taste [and ethics] are uniform in human nature: where men vary in their judgments, some defect or perversion in the faculties may commonly be remarked; proceeding either from prejudice, from want of practice, or want of delicacy....”⁵¹ That there is no ‘objective truth’ about these matters can and should be ignored, except insofar as it usefully militates against zealotry.

Indeed, I suspect that the views outlined in this paper have implications vis-à-vis our substantive moral judgments. We should, perhaps, jettison such terms as ‘right’ and ‘wrong’, in that they sound more objective than words like ‘good’ and ‘bad’. Nothing is [absolutely] *right* or *wrong*; it is merely good or bad (relative to something else), and there are gradations in its goodness or badness. Moral absolutism can seem lofty and noble, but it is simpleminded. Objective-sounding norms like ‘justice’ are, as Hume said, artificial; the real substance of morality consists in its prescriptions for well-being and against suffering. These are, by and large, the only ethical norms we ought to take seriously, given the comparative unpretentiousness of their pretensions to objectivity, as well as their immediately concrete implications and their direct relation to universal intuitions. Abstract principles, such as Rawlsian justice and Nozickian liberalism, have an air of ‘objectivity’ about them, as if they are justified by pure reason—which is impossible. Every such justification presupposes values and is merely an elaboration of values. Why not, then, forego the involute philosophizing, the building castles-in-the-air, in favor of alleviating suffering and promoting well-being on a case-by-case basis? Let’s forget about meaningless abstract principles like “Life for its own sake!” or “Marriage is a timeless institution between blah blah blah”: let’s be pro-choice and pro-gay marriage, because these positions are not self-delusively ‘objective’ or ‘abstract’ but instead promote concrete well-being, which is all that matters and is the only real substance to morality. And in any case, let’s abandon the enterprise of grand ethical system-building, which gets us nowhere.

A complete elaboration of the foregoing thoughts would occupy several books, but I hope I have at least sketched a few solutions to some of the perennial problems of moral philosophy.

⁵¹ Hume, *op. cit.*, p. 265.

The main point of this paper has been to show that the dichotomy between 'objectivism' and 'relativism' is false; a third way is available.