

CHAPTER 10

The Metaphysics of Ethics

Earl Conee

Right and Wrong

Custodian

Tom was disappointed. He did not get the job he wanted. Custodian Cathy saw how he felt about it. Cathy made a special effort to give Tom a positive perspective. Cathy acted out of sympathy. She had no expectation of getting anything in return. Her consoling efforts brightened Tom's otherwise depressing day.

That was a kind thing for Cathy to do. It was morally right.

Coach

Mort was an unpopular insecure student. During a gym class Coach Curt ridiculed Mort's basketball dribbling. Mort felt humiliated. Curt did that just to get some cheap laughs from his players who were watching in the bleachers.

That was a cruel thing for Curt to do. It was morally wrong.

We are expected to endorse moral judgments like those. But don't they just parrot the conventional line? It is hard to see how there could be any real facts here. Aren't moral judgments just subjective?

That is not so clear. In fact it may be impossible for there *not* to be moral facts. Suppose that Roger believes that eating meat is at least sometimes morally permissible and Ralph believes that eating meat is never morally permissible. Those alternatives seem to be exhaustive; they seem to include all possibilities. But some possibility has to hold. So it looks as though one of those beliefs is *true*. A true moral belief is a *factual* moral belief; that is, a belief in a moral fact. Now it looks as though there have to be moral facts.

Here is something else that counts in favor of moral facts. If there are none of them, then the best-grounded moral evaluations are strangely mistaken. In the *Custodian* example it turns out that Cathy *didn't* do anything right, and in the *Coach* example it turns out that Curt *didn't* do anything wrong. What could prevent these evaluations from being true? It was definitely considerate and helpful of Cathy to do what she did. Why isn't that enough to make it right? It was definitely cruel and petty of Coach to do what he did. Why isn't that enough to make it wrong?

It is difficult to maintain that grounds like those are inadequate, *if anything is truly right or wrong*. But maybe morality is just a socially enforced pretense and it doesn't describe anything in the real world. What could make a moral evaluation objectively correct? If we think our moral judgments are genuine truths about people and their deeds, we'd better have a good answer to that question.

While the nature of moral reality may be obscure to us, we can note that the subjectivist side has its own trouble. For one thing, it is likely that denying moral facts is not even believable, when we take it seriously. No doubt we can fake it. We can say to ourselves, 'Nothing is really right or wrong.' But when it matters to us, we continue to believe in objective right and wrong.

To convince ourselves of that, all we have to do is to recall a time when someone who we love was treated meanly. We can't help but affirm, 'That was wrong.' That seems as true to us as any other fact. At least that's how it is for me, and I'll bet the same goes for you.

In the end, our resistance to giving up our belief in morality might be our problem, though. Maybe the resistance is a product of habit: we acquired moral values early, we were often encouraged to evaluate morally, and the tendency to moralize lingers on in us. But now we have no rational defense of it. If so, then our continued belief in morality gives no good reason to think that there are moral facts. Before accepting that conclusion, though, we can look for a way to defend our nearly irresistible thought that moral evaluations are sometimes correct.

Here is a basic metaphysical question about these moral judgments. What is it to *be* right or wrong? Is there an aspect of reality that some moral evaluations correctly describe? If so, what is it? If not, what are we talking about when we moralize?

Realism

We seek the substance of moral truths. We can try applying to moral judgments an attractive idea about truth in general. The idea is easy to understand. A judgment that places something in some classification is true when, and only when, the judgment is backed up by some *way that the classified thing is*. For example, a result of a car accident is correctly classified as a dent because one effect of the accident is a new way that part of the car is shaped: it is dented. A result of the accident is correctly classified as a headache because one effect of the accident is a new psychological way that someone is feeling: someone's head is hurting. These ways of things are properties that the things have. To make moral judgments is to attribute

moral properties.¹ In the *Custodian* example, Cathy's act of consoling Tom has the property of being morally right. That is why it is true to say that the act is right. In the *Coach* example, Curt's act of ridiculing Mort has the property of being morally wrong. That is why it is true to say that the act is wrong. Generally, moral evaluations are about real properties that some actions and people really do have. This metaphysical understanding of morality is known in philosophical circles as **moral realism** because it claims that there are real moral facts and properties. Is it also realism because it harbors no illusions? We'll see.

Real Trouble

What are these alleged moral properties? They are not like the familiar properties of the natural world. Nature contains quasars, quarks, quakes, and queasiness. Evaluations of acts, such as being right and being wrong, don't belong on that list even after we expand it well beyond the 'q's. Moral evaluations are nothing like physical properties such as mass and charge, or biological ones such as being alive and having a gall bladder, or psychological ones such as hating the taste of cilantro and wanting to have a nap. They are not measurable conditions that things can be in. The following report by a medical technician in a white coat couldn't be true: 'We were monitoring Coach Curt's conduct with our morality gauge while he was making fun of Mort. Sure enough, the ridicule registered on the morally-wrong region of the scale.' Even the wearing of an official white coat wouldn't make that report credible. There couldn't be any 'morality gauge'. Right and wrong are not quantifiable qualities that could be detected by an instrument.

¹ For more about this see Chapter 8, 'Universals'.

Naturalism

Maybe we went too fast just now. Maybe it was hasty to dismiss *all* of the psychological properties as things that are quite separate from moral evaluations. Some emotional reactions are quite closely related to making moral judgments. When we think about Cathy's considerate good deed in *Custodian*, we feel approval. When we consider Curt's ridiculing in *Coach*, it repels us. In general, we respond with approval to morally positive things and with disapproval to morally negative things. Maybe something about feelings like these is *definitive* of the moral evaluations.

Moral and immoral acts don't have to get emotional responses from anyone in order to exist. Maybe right and wrong come from *tendencies* of conduct to elicit feelings such as approval and disapproval. In other words, maybe a tendency to prompt some favorable feeling is the true nature of being morally right and a tendency to prompt some unfavorable feeling is the true nature of being morally wrong.

This sort of view applies **naturalism**. It 'naturalizes' right and wrong. That is, this identification answers the question of where in the natural world are the conditions that make moral classifications correct. It tells us that these conditions are partly in the acts and partly in us. They are tendencies to bring about certain emotional responses. Tendencies like these are just as much parts of nature as the tendency of sugar to produce sweet flavor sensations when we taste it and the tendency of an extremely loud noise to produce discomfort when we hear it.

Natural Trouble

This naturalistic approach has a tendency of its own. It strongly tends to provoke objections. A couple of the objections are philo-

sophical classics. Before we get to them, here are a couple of preliminary problems that make life difficult for this naturalism.

First, whose feelings count? Does every last person who considers an act have to feel the specified way about it? If so, that just isn't going to happen. Misanthropes aren't going to feel favorably toward some of the best things that people do. Sociopaths aren't all going to disapprove of some of the worst things that they themselves do.

This might seem to be trivial trouble. Why aren't most people's responses enough, leaving aside these outliers?

Deference to a majority would make trouble. The trouble comes from the sort of people who could have been in the majority. The naturalism is supposed to tell us what right and wrong really are. If it does, then it covers all of the possibilities. Yet the population of 'most people' could have been rigged with malign intent. A fiendish genetic modifier could alter the genetic make-up of the future population. Suppose that these altered people end up being most of the people who ever exist. Their alterations have them approving of the damndest things. If the feelings of most people are what counts, then the theory implies that the feelings of that malicious majority would make the damndest things right. That's wrong.

An Ideal Solution?

A fix for problems like this has been proposed. We can call someone an **ideal observer** if the person is optimally equipped to respond definitively. We can try to explain moral evaluations in terms of the feelings, pro or con, by which an ideal observer would respond, if an ideal observer existed. But what in the natural world is 'ideal' about an observer? The naturalistic hope is that some psychological properties will do the job.

The feelings of an ideal observer have to be perfectly informed if they are really to define morality. No factual omissions or

mistakes can be made. To insure this, an ideal observer must be all knowing. All bias and self-interested favoring must be excluded. To insure this, an ideal observer must have a perfectly impartial perspective.

An ideal observer needs some further psychological features. Being all knowing and impartial does not guarantee having any feelings at all. According to the theory, it is an ideal observer's feelings that make things morally good or bad, right or wrong. So if an ideal observer felt nothing, then the theory would count nothing as good or bad, right or wrong. That would be giving up on morality, not showing us what its place in nature is. If the ideal observer had feelings but was emotionally troubled, then the feelings would be distorted rather than definitive. Also, emotive idiosyncrasies cannot count. Tastes are emotive attitudes. But the theory must exclude all such irrelevancies. An ideal observer could not establish genuine moral values by being a cauliflower connoisseur or a hazelnut hater. Any such accidents of taste couldn't be morally definitive.

What is the rest of the best emotional constitution for an ideal observer? It complicates the problem that naturalists have to avoid using any moral evaluations in specifying the psychological requirements. It would help a lot, for the purpose of getting the emotive responses that naturalists seek, to require an ideal observer to be 'virtuous' or 'ethically sensitive'. But that sort of requirement relies on ethical evaluations rather than explaining them in terms of natural properties.

It is not at all clear that a naturalistic specification of an ideal observer can require just the right feelings. And even if it can, the classic objections that we are about to discuss would still apply. So, having seen some hope of solving the preliminary problems for naturalism, let's proceed to those objections. It'll be harmless to pose the objections against a simple version of the general naturalist idea:

Prompting Attitudes (PA): For any act, A, to be morally right is for A to have the tendency to produce approval in those who consider it; for any act, B, to be morally wrong is for B to have the tendency to produce disapproval in those who consider it.

Classic Complaints Against Naturalism

PA says that being morally right is just the same thing as having the tendency to prompt approval in those who consider the act. The relation of rightness to the tendency is supposed to be a numerical identity.² The two classic objections to PA seize on this fact.

First Classic Complaint: Euthryphro

Recall Cathy in the *Custodian* example. She generously takes the trouble to console Tom. When we consider this act, we respond with approval. What is it about the consoling that prompts this reaction? Well, it was considerate and selfless. But that's only part of our reason. Someone could have been considerate and selfless in the service of knowingly aiding the murderous efforts of a cold-blooded killer. That would not have met with our approval. So there's more involved in getting our approval. The objection contends that part of the whole story about what secures our approval is that we think of the consoling as *the right thing* for Cathy to do. In fact, the critic contends, the rightness is the clincher for gaining our approval. This sets up the crucial point.

Suppose that thinking about the property of being right does help to prompt our considered approval of the act. Nevertheless, when we consider the act carefully and find that we approve of it, we don't think about how other people feel. We think about Cathy, Tom, and the consoling. The feelings of others are seldom,

² For more about this see the discussion of 'numerical sameness' in Chapter I.

if ever, on our minds. So the property in PA of *being something that provokes approval in all upon consideration* does not play a role in that thinking. Yet the property of being right is on our minds. That is a difference—a difference in which property we are thinking about. That’s enough. *Any* difference excludes their numerical identity. Hence, PA incorrectly asserts that these are one and the same property.

This is a version of what some philosophers call a ‘Euthyphro problem’. Plato’s dialogue *Euthyphro* contains the original rendition of this sort of objection.

The same problem affects other moral evaluations identified in natural terms. Here is a naturalistic identification about being good that has been used to illustrate the problem. It identifies the good with what we want ourselves to want.

Desire to Desire (DD): For something to be good is identical to its being something that all desire to desire.

Happiness is good. According to DD, another way to state the same fact is that we all desire to desire happiness. Suppose we do. Why do we want that? Well, we know that happiness is often a pleasant condition. But we also know that a state of happiness sometimes isn’t actively pleasant. Having happiness isn’t being on a perpetual high. The pleasure of happiness isn’t the whole story about why we want to want happiness. Eventually it becomes clear that at least many of us who want happiness want it because we appreciate this point: all in all it is a good thing to be happy. In this way the goodness of happiness is part of what gets us to desire to desire it. But the property of *being something all desire to desire* is *not* getting us to want it. We *do* think about the goodness of happiness and we *don’t* think about how widespread any desire for it is. So again the properties differ. The property of being good played a psychological role that the property of being something all desire to desire did not play. Therefore, they cannot be identical.

The reasoning amounts to this. Suppose I want to want happiness, and I think about why that is. I think to myself, 'Because *it's good*'. I don't think to myself, 'Because *we all want to want it*'. Therefore, being good isn't the same thing as being what we all want to want. We have reached the conclusion that DD is untrue.

Second Classic Complaint: Open Questions

The other traditional objection to naturalistic identifications like PA and DD is called the Open Question argument. The twentieth-century philosopher G. E. Moore devised it. DD was one of Moore's targets.

Here are two questions:

Q1: Is each good thing something good?

Q2: Is each good thing something all desire to desire?

DD tells us that there is just one property that Q1 and Q2 are asking about, namely goodness, and they are asking the same thing about that property. If that is so, then clearly in English the property is invoked by the word 'good' and equally it is invoked by the phrase 'something all desire to desire'. They bring to mind the very same property in those who understand their meaning.

At this point the Open Question argument takes a seemingly modest step. From the observation that DD implies that Q1 and Q2 ask the same thing about the same property, the Open Question argument infers that given DD, there is *a mere difference in wording* between the question that is asked by Q1 and the question asked by Q2. In their substance, Q1 and Q2 ask the same thing.

The argument continues. When we step back from what DD alleges, it is clear to us that what Q1 asks is *strikingly different* from what is asked by Q2. Q1 is an idle question. Of course, each good thing is a good thing. There is no room for reasonably

wondering about that. We can mark this rational emptiness of Q₁ by calling it a 'closed question'.

Just as Q₁ is closed to reasonable doubts, Q₂ is open to them. It can be entirely sensible to wonder whether every good thing is something that everyone wants to want. Maybe not everyone has happened to think of all good things, much less have some desire about each of them. It is at least reasonable to wonder about that. Any such doubt about the answer to Q₂ shows that Q₂ could be a rational question to ask. It is an 'open question'.

So as a matter of fact Q₁ and Q₂ do differ in substance, not just wording. They differ in whether or not it is reasonable to wonder about the answer. So Q₁ and Q₂ must *not* ask the same question. The Open Question argument contends that this is where DD goes wrong. The argument has inferred from DD that Q₁ and Q₂ ask the same question and they differ only in wording. But we've just seen that they ask different questions: Q₁ is closed and Q₂ is open. The argument concludes that because DD has this erroneous implication, DD is untrue.

DD can be improved. A careful limit can be placed on the people who DD requires to want to want something in order for it to be good. One improved theory restricts these people to those who have thought about a full range of topics, so that they haven't overlooked anything good.

The same sort of Open Question argument applies against the improved account. It remains reasonable to wonder whether even the improved condition really does succeed in isolating exactly the good things. We can easily wonder: couldn't there be sensible reasons why people might not want to want good things like happiness? And anyway, why does everyone have to have any wants about wants? Maybe some people don't want to want anything, just because they have never thought about that. Also, don't Buddhists want *not* to want anything? Maybe in spite of

themselves they also do want to want all of the good things . . . but why would they *have to* have his inner conflict? We are now sensibly wondering about the answer to Q2. Meanwhile, the closed Q1 remains as un-wonder-about-able as ever. The Open Question argument infers from this difference in the questions that goodness is misidentified by DD.

How could any naturalistic identification of goodness avoid this? How could it *not* introduce some new idea, an idea of something natural, an idea that is not just contained in the idea of being good? If they all do that, then they all make themselves vulnerable to Open Question arguments.

More on the Classic Complaints: A Misgiving

You might think that both of these complaints amount to nit-picking that doesn't threaten anything except your patience. You might think that pursuing verbal fine points about phrasings cannot show us anything about the nature of major ethical evaluations like right and wrong, good and bad. It might seem that this criticizing is just playing with words.

We should give these objections a fair hearing, though. It is in our own interest. Progress in metaphysics is difficult. We have to take full advantage of whatever intellectual resources we have. We know some things about words and their meanings, and we know some things about our own attitudes. The Open Question argument exploits this. The reasoning can seem petty. But it also seems crafty. It does have some apparent force. It appears that our verbal knowledge has a chance of getting us somewhere concerning the nature of moral evaluations. We've got to take seriously any reasoning that might turn out to make progress. If it does succeed, then we're getting some metaphysical work done with words and not just playing with them.

Non-Nature

Before looking more critically at the Open Question argument, let's think about what its total success would imply. Let's suppose for now that the reasoning shows that moral evaluations are not about any *natural* properties. Maybe they are about some other properties. How can we be sure that nature is the whole of reality? Many moral judgments do seem true. For example, it still seems quite clear that conduct like Cathy's in *Custodian* is right and conduct like Curt's in *Coach* is wrong. For now, let's stay with moral realism. So what makes moral judgments like these correct includes the existence of the properties of being right and being wrong. The new thought is that these are real properties that such acts have, but *non-natural* properties. That's what G. E. Moore thought. What's wrong with that idea?

Natural Dependency

Nothing is conclusively wrong with it. But it faces difficulties. One of them arises from something that is grandly called 'supervenience'. Fortunately, the idea behind the term is interesting and readily understandable.

The supervenience is something that we take for granted. Suppose that we think about Sidney. We know Sidney has a wonderful disposition. She is considerate, generous, brave, and honest. She is never intentionally harmful in any way. She has all of these good psychological characteristics and no negative ones. This tells us that Sidney is a good person. Suppose that in some distant part of the universe there is a Duplicate Earth that shares every natural feature with Earth. Duplicate Sidney is there. She is just like our Sidney in all natural characteristics. This tells us that Duplicate Sidney too is a good person. She must be, because she is the very same sort of person as Sidney is, in all the ways that matter for

being good. It would be *loopy* to think that Sidney is good and Duplicate Sidney is just like her, except that she is not good. Now here's the point about moral supervenience. When we see that it would be ridiculous to count as good our Sidney but not Duplicate Sidney, we are relying on the natural characteristics of our Sidney to *determine* that she is good. In philosopher's jargon, we are relying on the moral status of being a good person to *supervene on* the person's natural features.

Whenever we find a moral difference, we take it to derive from some natural difference in the psychology of the people involved or their physical or social circumstances. This act was wrong, and that one was not, because this one had malicious intent and that one didn't, or this one was damaging while that one was harmless. It's always something like that. We count on the existence of some natural difference to induce the moral difference, because we see natural conditions of some sort as determining what the moral situation is. This determining by the conditions of people and their circumstances of the moral status of their acts is *not* cause and effect. It is more inevitable than that. The natural features completely settle the moral status, no matter what. It looks as though the same goes for all examples of *any* moral status: good or bad, right or wrong, permissible or forbidden, and so forth. That is:

Moral Supervenience (MS): In any possible case of a moral evaluation, there are some natural conditions that are necessarily sufficient for the evaluation to apply.

It is worth pausing to think about whether we can come up with an exception to MS. MS asserts that a moral status is always determined by natural conditions. An exception to MS would be like this: two possible examples are exactly alike in every aspect that is part of the natural world—physically, socially, psychologically, and every other natural way. Yet something has some moral status in the one case—an act is morally right, someone is a

moral person, or the like—while in the natural duplicate case that something is *not* morally the same. Could that happen? . . . On reflection, we don't find anything that could make the natural duplicates morally different. It seems impossible. No exception leaps to mind. On further reflection, no exception lumbers to mind.

MS is looking true. That's interesting. (It is!) And it will make a difference later, when we get to the view known as emotivism. What it does for us now is to prepare us for a criticism of non-naturalism—the view we are considering that says that moral properties are non-natural properties.

MS does not directly comment on non-naturalism. But MS shows us that non-naturalism creates a mystery. MS tells us that moral evaluations are necessarily settled by natural conditions. Yet suppose that non-naturalism were correct and no moral property was identical to any natural one. If that were true, then why would natural conditions always *determine conclusively* whether or not the moral property was there? Why couldn't the moral property get detached from its natural underpinnings? Moral properties aren't glued onto natural ones. Nothing like gluing would help us to understand the connection anyway. Even the best glue doesn't necessarily hold. In contrast, MS tells us that the natural-to-moral link is absolutely necessary. Non-naturalism leaves this as a total mystery. That is a liability.

Long Odds

As *Star Trek* fans know, the Borg is an immensely powerful collective of ruthless, relentless, ceaselessly adaptive invaders who are bent on universal conquest. Their well-justified slogan: 'Resistance is futile'. Science is the Borg of inquiry. It has been overwhelmingly successful at eventually explaining things, and it keeps getting better and better. It looks futile to resist the

conclusion that moral properties fall within the realm of science, if moral properties exist at all. Nothing we know goes as far as to establish the impossibility of a non-natural realm of properties, out of the reach of science. But given the success of science, its existence would be a bad bet.

Back to a Classic Complaint: Questioning Open Questions

We should re-examine the reasoning that gives credibility to non-naturalism; namely, the Open Question argument. At a crucial juncture the reasoning makes a dubious inference. It starts from the safe thought that, given the naturalist identification DD, the word 'good' and the phrase 'something all desire to desire' brings to mind the same property. But the argument then takes a fateful step. It infers that according to DD, the phrase 'something that all desire to desire' is *just a rewording* of 'good'. Considered carefully, that looks like a misstep. It seems to ignore a possibility: the two phrases might bring to mind the same property, but *in conceptually different ways*. The concept that the word 'good' places before our minds can differ from the concept given by the phrase 'something all desire to desire', even if they are concepts of the same property.

This becomes a highly credible possibility once it's raised. The same sort of thing pretty clearly does happen when we have both non-scientific and scientific concepts of kinds of substances. Rubies are popularly thought of as being a certain type of red gemstone. Rubies are less often thought of as being red crystalline aluminum oxide with trace iron. But that's what it is to be a ruby. Science has discovered that being a ruby is having that molecular composition. When people first conceived of the gems as rubies, it would have been a completely 'open question' whether that molecular composition was the nature of a ruby, if anyone had happened to think of that composition at all. But

the openness of that question cannot show that the science is mistaken. It just shows a difference between pre-scientific and scientific conceptions of what rubies are. There is no apparent reason why the same couldn't go for moral properties—different concepts stand for same moral property.

This is powerful evidence that the Open Question argument makes an invalid inference. The argument makes an inference from safe premises to the conclusion that given DD, the two questions about goodness, Q₁ and Q₂, have to be only *verbally* different. But we've seen that the questions could also be *conceptually* different. They could bring to mind different concepts of goodness, even if they were both concepts of the same property. That difference might be what opens up question Q₂ while question Q₁ is closed. If so, then the Open Question argument doesn't refute the identification of a moral property with a natural one.

New Naturalism

In the later part of the twentieth century, new ethical naturalists made use of this sort of rebuttal to Open Question arguments against naturalism. The new naturalists also applied another good idea. The idea is that some terms apply to something because they have the right *causal* link to it.

Names are prime candidates. For instance, suppose we have a friend named 'Mark'. The causal view is that Mark can have that name now because someone said, while pointing to newborn little Mark, 'We're calling him "Mark"'. When we now use 'Mark' to refer to our friend Mark, a series of past uses of the name goes back, by cause and effect, tracing through other minds and other mouths, to that first linking of the name to him. That connection makes all of those uses of the name refer to our friend Mark.

Here is a puzzle about names. Mostly we use the name 'Mark' for our friend Mark. Mostly we are thinking and talking only about our Mark. But lots of guys are named 'Mark'. So how do our uses of 'Mark' find their mark?

The causal view has a simple solution. A guy named 'Mark' gets referred to with the name when a particular use of the name causally traces in the right way back to that guy alone. Our 'Mark' traces to our friend Mark only. Problem solved. Score one for the causal view.

Terms for general kinds of things can operate like that too. It can be that an ark is called an 'ark' because the term was given to that kind of ship by pointing to one of them and saying something like this: 'Let's use "ark" for things like that.' This would be coining a term for arks by causally linking the word to some property that the example has—the property that makes it a 'thing like that'.

The new naturalists think that terms for ethical kinds, terms such as 'good', 'bad', 'right', and 'wrong', refer in this causal way. Being naturalists, they think that ethical terms designate things in natural kinds. Each ethical term applies to whatever has the natural property to which the term has been linked by some proper causal connecting.

Which natural properties are the moral ones? If the new naturalism is correct, then we cannot figure that out by just thinking about what we mean by the ethical terms. That thinking doesn't enable us to track down the properties that are at the beginnings of the causal series. Only investigating the causal lineage could decide conclusively the nature of the natural properties. We can make guesstimates. We have the close association between moral evaluations and feelings. It can guide us. One simple guesstimate is that 'good' is linked to happiness, 'bad' is linked to unhappiness, 'right' is linked to promoting happiness, and 'wrong' is linked to promoting unhappiness. Or maybe 'good' and 'right' are linked to kinds of things that we feel

favorably toward; 'bad' and 'wrong' are linked to kinds of things that we oppose. These are among the sensible conjectures.

In any case, the natures of the linked properties are not at all obvious from our knowing the meanings of the ethical terms. So the new naturalism lets it be easy for us to make big mistakes about what really has the ethical properties, just as we can make big mistakes about what is a real ruby. And that seems quite true—people do sometimes make big mistakes about what is moral. You'll have your own favorite examples. They might not match mine. If so, then one of us is making one of those big mistakes. Providing this basis for the fallibility of our moral judgments is an asset of the new naturalism.

Trouble in the Twin Cities

The new naturalism seems very promising. So, you may ask—having noticed that philosophers apparently have objections to all philosophy—what do philosophers have against it? Well, the view suffers from a certain detachment.

The complaint can be brought out by a tale of two cities (a tale derived from one told by the philosophers Terry Horgan and Mark Timmons). The cities are much alike. The residents of each city speak a language that sounds exactly like English. A key social difference exists. In City One, happiness has a more central role in people's lives than it does in City Two. In City Two, giving and getting respect looms larger in people's lives than it does in City One. This social difference is just enough to have the following consequence. As residents of City One use the term 'right', it is causally connected so as to apply to acts that promote happiness. As residents of City Two use 'right', it is causally connected so as to apply to acts that attract respect.

Now suppose that a Resident Of City One, **Roco**, is discussing a certain scandalous act with a Resident Of City Two, **Roct**. (The act will not be further specified, to avoid needless wallowing.)

Roco and Roct both know that the act promoted happiness and attracted no respect. Roco says, 'Like it or not, that act was right.' Roct replies, 'So you say. I say that it was not right.'

It seems clear that Roco and Roct disagree. We are in agreement about that. But then the new naturalism is wrong. The new naturalism tells us that what Roco correctly calls 'right' is anything that promotes happiness. It tells us that what Roct correctly denies to be 'right' is anything that fails to attract respect. The scandalous act does promote happiness and does not attract respect. So by the new naturalism, both are telling the truth. According to the new naturalism, then, Roco and Roct are just telling different truths, not disagreeing. They *are* disagreeing, though—we agreed about that near the beginning of this very paragraph. Since the new naturalism wrongly implies a lack of disagreement in the exchange between Roco and Roct, it is untrue.

Troubling Emotional Involvement

Stepping back from this specific objection, we can see a general problem for moral realism. It tells us that moral evaluations attribute properties to their subjects, just as ordinary descriptions do. The general problem is this. Moral evaluations are more intimately entangled than that with feelings, intentions, and advice. For example, we might describe a nose punching as 'injurious' in order to make a formal report of the fact that the punch injured the punched. That's a description that we could make with total indifference to what was done. But we wouldn't call the punch 'morally wrong' unless we cared about it in some way. We would have some negative attitude that would be engaged by thinking about the punch. It would get us to rate the act as wrong.

With this background in view, we can see that when Roco counts an act as 'right', we take it for granted that Roco's

emotive attitudes are involved and they are in some way favorable to the act. When Roco responds by counting the deed as 'not right', he is registering that he is not on board with Roco about it—he does not feel favorably in that way. Their verbal exchange gives voice to a conflict of emotive attitudes. That would not be so, if they were just reporting on whether or not the act had some property. When we think of the exchange between Roco and Roco as being a 'disagreement', it may well be that some such conflict of attitudes is what we are discerning. Moral realism makes no place for that in its interpretation of what they are saying.

Emotivism

There are benefits to taking to a philosophical extreme the role of emotions in moralizing. The extreme idea is that moral evaluations are verbal outpourings of emotive attitudes. The evaluations have nothing to do with moral properties. The **emotivist** view says that there aren't any of those. So it denies moral realism.

Here is one asset of emotivism. Since moral properties don't have to exist in order for us to make moral evaluations, a problem that we've seen for moral realism is gone. Moral properties don't have to show up anywhere, either in nature or in some non-natural realm. That's good because we were having trouble finding any properties that seemed to be fully qualified for the job.

Simple Emotivism

When we morally evaluate, exactly what emotional thing are we doing according to emotivism? A radical thought is that we are not saying anything true or false, we are just giving vent to

emotions. Some emotivists hold that when we call Curt's ridiculing of Mort in the *Coach* case 'wrong', for example, we are using the word to give verbal release to a negative sentiment toward the ridicule. We have some other terminology that uncontroversially does that sort of thing, words like 'boo' and 'eww'. That language can be recruited to illustrate this radical emotivism. It holds that what we mean by calling Curt's ridicule 'wrong' could be revealingly reworded like this: 'Curt's ridicule—hiss!' Similarly, our saying that Cathy's consoling in the *Custodian* case was 'right' is revealingly reworded like this: 'Cathy's consoling—hooray!' We aren't classifying the conduct by attributing a property to it. We aren't saying anything about how any part of the world is. What our moralizing does is to give vent to our emotions.

Expressive Enhancements

This is the simplest version of the emotivist approach. Improvements exist. First, when people moralize, they aren't often feeling the crude aversions and attractions that are expressed by hissing or cheering. Emotivism isn't limited to relying on any such simple feelings. Emotivists can say that we use moral language to express certain serious and careful forms of approval and disapproval.

With that elaboration, emotivism can cope with unenthusiastic moralizing. Suppose that you tell me: 'For Barney's own good, I morally ought to clue him in about his poor singing abilities, even though this will be painful all around.' You wouldn't be feeling anything like cheering when you made this positive evaluation of hurting Barney's feelings for his own good. Emotivists can agree. They can cite subtler positive or negative sentiments. In this case what you'd be feeling toward telling Barney about his singing would be some regretful sort of favoring.

A second improvement results from emotivists saying more about what gets us to make moral evaluations. This will help their view to accommodate something important that we've already noted: the highly plausible claim of MS that the moral derives from the natural. At first, MS looks bad for emotivism. Emotional reactions can be irrational. Nothing guarantees that every last natural duplicate of something that we moralize about will get us to feel the same way. Maybe we like some of the duplicates and dislike others, just on a whim. Emotivism tells us that without the same sort of feeling we wouldn't make the same moral evaluation. MS implies that things have got to be *the same* morally whenever they are the same naturally. Cases of irrationally differing reactions to natural duplicates seem perfectly possible, even likely, for the whimsical likes of us. Doesn't MS tell us that the moral differentiations that emotivism finds here are mistakes made by emotivism?

They are not mistakes if emotivism is supplemented as follows. Moral evaluations, when they are made sincerely and with full understanding, are made *on a certain basis*. When we sincerely and thoughtfully call something 'right' or 'wrong', this is because we take it to have certain natural properties, and they get us to have certain emotive attitudes. For instance, in the *Custodian* case what gets us to admire the conduct would be something like Cathy's considerate thoughtfulness in her consoling of Tom. In the *Coach* case what we deplore would be something like Curt's callousness and his cheap attempt at ingratiation in his ridiculing of Mort. We have some such natural properties as bases for our moral attitudes. Relying on these bases makes us all set to have the same attitude toward whatever we think has the same natural properties. So no wonder MS is so plausible. Our careful reflections don't turn up any examples that appear to violate MS, because in careful uses of moral terms the same natural basis gives us the same emotive reactions. According to

the supplemented emotivism, it is this sort of reaction that we are voicing in making a thoughtful moral evaluation.

An improved version of emotivism like this is known as **expressivism**. The view is now looking pretty good. As you no doubt expect, however, some philosophers think that expressivism has some discrediting liabilities. Here are two of them.

Conditional Trouble

First, when we ethically evaluate we don't just engage in isolated evaluative outbursts. Sometimes we reason about right and wrong using more complicated claims. Emotive expressions don't seem rational enough for this task. Here's an illustrative piece of reasoning, with rational defenses of its premises in parentheses:

Deception Reasoning

Premise 1: Some intentionally deceptive lies are morally permissible (such as lies that harmlessly spare someone from great distress who is about to die).

Premise 2: If intentionally deceptive lying is sometimes morally permissible, then so are some intentional deceptions that are evasions but not lies (since intentional deception is the worst aspect of the lying and evasions are otherwise no worse).

Conclusion: Some evasions are morally permissible.

Deception Reasoning defends its modest conclusion pretty well. Expressivists must have some account of what is reasonable about it. All ethical reasoning poses a challenge to expressivism, just because it is reasoning. If moral evaluations are emotive expressions, not true or false claims about how things really are, then how can we *reason* about them? Isn't the point of reasoning to derive *truths* from *truths*? Expressivists owe us some explanation of truthless rationality.

Expressivists do have an explanation to offer. They propose that our reasoning about morality brings out our derivative emotional commitments. This is an emotional parallel to the standard view that reasoning about truth brings out derivative truth. Having noted this issue, let's set it aside and focus on a further problem that Deception Reasoning brings up.

The further problem is to explain what is meant by some compound moral sentences such as the second premise, P₂. Expressivists can credibly say that a simple moral claim, such as the one made by premise P₁, expresses some attitude like this: toleration toward some intentionally deceptive lies. But what emotionally expressive job is done by a conditional claim such as premise P₂?

Notice that someone who sincerely affirms P₂ need not feel any particular way about intentionally deceptive lying. For instance, Sasha *reviles* those lies because they are instances of what she regards as the loathsome practice of intentional deception. Still, she'd concede that some evasions *would be* sometimes okay *if* that sort of lying was okay. On that basis she affirms the conditional claim P₂. Sylvester affirms P₂ because he feels *positively* toward all lies. He thinks that they pose challenges that toughen us up in our intellectual lives. Sylvester affirms P₂ when he notes the consequence that some evasions are helpful in this way too. Sasha and Sylvester seem to be basing their affirmations of P₂ on drastically different attitudes. It looks as though no one emotive attitude could give P₂ its meaning.

Expressivists do have proposals about what the attitude is. One leading idea is that affirming P₂ is voicing a complex emotive attitude, something like this: opposition to the combination of a tolerance of some lies and an intolerance of all intentional deceptions. Those like Sasha who affirm P₂, while reviling all intentionally deceptive lies, should be prepared to unite in having that complicated attitude with those like Sylvester who affirm P₂, while liking all lies. Their feelings make them alike in this

way: they are committed to opposing the combined attitudes of tolerance for intentionally deceptive lies and intolerance for intentional deception.

That solution isn't irresistible. Maybe Sasha and Sylvester should be *prepared* to share that complicated attitude of opposition, because by affirming P2 they are somehow *committed* to it. But is that negative attitude toward that combination of attitudes *actually present* in those of us who affirm P2? I am prepared to confess that the attitude didn't seem at all familiar to me, even after I convinced myself to affirm P2. It is doubtful that we who affirm P2 all have any such elaborate attitude toward attitudes. If not, then the proposed expressivist interpretation of P2 is in trouble. Expressivism asserts that our moralizing serves to voice some attitudes that we have; it serves to release them verbally. We can't verbally release an attitude that we don't have.

Pondering

Another problem for expressivism derives from another mental role that can be played by moral claims. Suppose that we simply consider the claim that lying is sometimes morally right. In doing the considering, we seem only to be holding that claim before our minds. We are just calmly focusing on the allegation that it makes. What emotional attitude toward lying might be at work when we coolly contemplate the claim?

Expressivists can propose that this pondering is taking an attitude of quizzicality. It would be an attitude we can put in other words by saying something like this: 'Some cases of lying, hmm.' But that proposal still seems too emotional. In order to consider the claim that lying is sometimes right, we needn't have any feelings stirred up at all, not even feeling quizzical about it. We can have it in mind without so much as a mild curiosity or any other sentiment about it. Cool contemplation just isn't emotionally engaged.

That goes against what expressivism tells us that moral claims do for us. Apparently we can bring to mind a moral claim without its serving any function involving our emotive attitudes.

Ethical Errors

We haven't found a fully satisfactory way to make good sense of moral properties, whether as natural properties or as non-natural properties. We haven't found a fully satisfactory understanding of moralizing without moral properties either. What's left? The last alternative that we'll consider is **error theory**. It says that morality is all a mistake. When we moralize, we are trying to tell the truth about how the world is. We make assertions about how things are morally. But no moral properties exist to enable us to assert moral truths. We are always in error.

For instance, here is a moral claim that it is difficult not to believe:

Wrong to Agonize Innocent People for No Reason (WAIPNR): It is morally wrong to subject innocent people to agony for no reason.

Despite the credibility of WAIPNR, error theory implies that it is untrue. *Nothing* is morally wrong. Claims about acts being wrong attribute a moral property, and there isn't any such property for the claims to tell the truth about (we'll see shortly why not).

Doesn't error theory render itself ridiculous right there? It seems to be affirming the *denial* of WAIPNR:

Not Wrong to Agonize Innocent People for No Reason (Not-WAIPNR): It is not morally wrong to subject innocent people to agony for no reason.

Yet affirming Not-WAIPNR seems preposterous. Out-and-out affirming that such terrible inexcusable conduct *isn't* wrong seems to be upholding a crazed morality.

One careful version of error theory does *not* affirm Not-WAIPNR. The version holds that all moral claims, both positive ones like WAIPNR and negative ones like Not-WAIPNR, *presuppose* the existence of moral properties. That is, all moral claims in effect allege that the properties exist, whatever else they say. WAIPNR tacitly says that moral wrongness exists while explicitly asserting that it characterizes agonizing innocents for no reason and Not-WAIPNR tacitly says the same thing while explicitly denying the same explicit assertion. As a result of their tacit allegations that moral properties exist, all moral claims are untrue. So both WAIPNR and Not-WAIPNR are untrue.

Another version of error theory *does* affirm Not-WAIPNR. The torture that it is about is not morally wrong, because nothing is. These error theorists urge those who doubt their view to be careful. If we hear Not-WAIPNR affirmed, we expect that something else is going on in the affirming person's mind too. We expect anyone to think that whatever is not wrong is permissible. So we expect anyone who affirms Not-WAIPNR also to hold that it is morally permissible to agonize innocents arbitrarily. Contrary to this expectation though, error theorists do not also hold those things. They deny them. Again, they say that nothing is morally wrong *and nothing is morally permissible*, because there are no such properties.

It also helps the plausibility of error theory to note that what the error theorists deny are specifically *moral* evaluations. This allows them to be consistent in variously vigorously opposing appalling conduct such as the arbitrary agonizing of innocents. They can find it repulsive. They can hate it. They can favor severely punishing it. They can be willing to die to prevent it. They just can't consistently *moralize* about it.

What convinces error theorists that moral properties don't exist? As they see things, all alleged moral properties have some fairytale-like aspect to them. For example, consider the alleged property of being morally obligatory. Error theorists think that for an act to have the property of being morally obligatory, the obliged person would have to be subject to some inescapable rule. It must be a rule that is built into the fabric of the universe that *demand*s the act, whether or not it is actually performed. Yet nothing in the universe makes this sort of demand. Maybe laws of nature 'demand obedience'. But they do this only in the sense that they do get followed, no matter what. Error theorists point out that any other sort of 'universal inescapable demand' is just a fantasy. They say that nevertheless, that is what it would take for some act to be morally obligatory. Error theorists conclude that there's no such thing as being morally obligatory.

Similarly, some error theorists contend that for there to be any such thing as the moral property of being good, the property would have to make whatever had it *intrinsically attractive*. It would be appealing to all, regardless of psychology and background. But nothing is *that* irresistible. Error theorists conclude that the property doesn't exist.

Error theory proclaims that there is no truth in morality. We might feel that we must oppose this view because it seems to legitimate any conduct at all, however horrendous. But error theory definitely doesn't *morally* legitimate any conduct. It does deny that conduct is ever morally objectionable. But it does not encourage us to be indifferent to whatever is done. Error theory allows well-founded favoring of some conduct and well-founded opposition to other conduct. Error theory is consistent with our having good grounds for these attitudes because the conduct matters to us in any of numerous non-moral ways—we enjoy it or it disgusts us, we get inspired by it or depressed by it, we love it or hate it, and so forth.

Anyway, even if error theory did have dangerous implications, it might be that error theory turns out to be the most reasonable view of the metaphysics of morality. Before we accept it though, we should give it some critical attention.

Errors About Errors

One liability of error theory derives from the extreme credibility of the likes of WAIPNR. We don't get to be rational in believing things like WAIPNR just because they strongly strike us as being true. We are altogether too fallible about the facts, even when a claim seems quite true.³ There can still be better reason to deny it. But if a claim has seemed as close to irresistibly right as has WAIPNR, for so long, to so many reasonable people, then considerable caution is warranted before we deny its truth. We run a great risk that some truth is there but it has been misinterpreted or it has been faulted for implications that it does not really have.

We should scrutinize how error theorists defend their astounding assertion of massive moral error. As we have seen, they assert something along these lines: alleged moral properties like being obligatory need unbelievable conditions to hold in order for the properties to exist, such as the existence of demands made by rules that are inherent in the universe. Plainly, no such rules exist.

Maybe moral properties seem unbelievable to error theorists because they exaggerate their requirements. For instance, does the existence of morally obligatory conduct truly depend on the existence of cosmic demands? Maybe saying that moral obligatory conduct is 'demanded' means only this: we *morally must* perform any morally obligatory acts. That is, it has to be that if

³ For some locally available examples, consider the strong credibility of certain jointly inconsistent thoughts about freedom and about universals, as we discuss in our chapters on those topics. Despite their credibility to us, at least one of each of those groups of highly credible thoughts must be untrue, because they conflict with one another.

we do not perform those acts, then any alternative that we do take is immoral. This involves no literal demands. It says only that we must fail morally if we do not do what is morally obligatory. There is nothing fantastic in that. The error theory view that morality is a mistake might be undercut by errors like that one.

Conclusion

The metaphysics of ethics is not easy. Although the prominent approaches show some promise, they all face trouble. We could get exasperated and give up. But that would be hasty. For one thing, investigating the metaphysics of ethics is mind-expanding. For instance, we have seen possibilities that are good to know about. We probably wouldn't otherwise have noticed that there is a way for an act to be neither morally permissible nor morally impermissible. In any event, questions of the reality of morality are intriguing and important to us. That makes it wrong for us to give up investigating them. And there we have another ethical evaluation to try to understand metaphysically.

FURTHER READING

As a next reading about the metaphysical issues discussed here, a very helpful resource is the Stanford Encyclopedia of Philosophy article 'Metaethics' by Geoffrey Sayre-McCord: <<http://plato.stanford.edu/entries/metaethics/>>. Its bibliography includes the classic books and articles, and much more.