

Allan Gibbard, "The Reasons of a Living Being" (2002)

When I came to the University of Michigan twenty-five years ago, Charles Stevenson had just retired, and I came to occupy his budget line. Over the next few years, this seems to have had a deep effect on me. Previously I had thought a lot about moral disputes and what's at issue in them, but I'd just been baffled; I hoped some solution would turn up. Now some Stevenson-like ways of tackling the puzzle began to occur to me, and I convinced myself that they have more power than I had previously thought possible. Perhaps I had good reason to be convinced - or perhaps you'll find it was just the subliminal influence of the budget line.

The puzzle about moral issues was Moore's puzzle, the one that G.E. Moore made especially vivid a century ago. As we all learn at our philosophical parents' knees, Moore argued that moral questions concern a non-natural property. When we try to settle a moral question, he maintained, we're not in the same line of inquiry as when we use empirical, scientific methods to inquire into the natural world. Notoriously, Moore had an "open question" argument which seems in retrospect to be fishy, and a "naturalistic fallacy" which he put in lots of different ways, all of which seem to beg the question. Later on he saw his first book as terribly confused. But he had another line of argument which, it seems to me, just won't go away; I call it the "What's at issue?" argument.

Jack, imagine, claims that all pleasure is good in itself, but Jill says that guilty pleasures are not in themselves good. So Jack says that all pleasure is intrinsically good, and Jill disagrees. What's at issue in all this? The two disagree about something, sure enough - but what? Jack, imagine, adds that after all, 'good' just means pleasant. But if he's right about what 'good' means, then they can't be disputing whether all pleasure is good, for they both agree that all pleasure is pleasant. Take any preferred definition of 'good', Moore argued, and we can construct a similar puzzle for it. Now whether any Moore-like argument can be made to work is still a matter of controversy, to be sure, but Moore does, with this argument, offer us a broad test for any account of what moral claims consist in. Ask what's at issue; that's the test. What, according to the account, is at issue in moral disputes? What does the disagreement consist in? Some accounts, even today, won't have plausible answers to this question. Charles Stevenson thought that Moore's arguments on this score worked, or at least that he could find arguments like Moore's that worked. And Stevenson had an answer to the "What's at issue?" challenge: Jack's for and Jill isn't. What's at issue is what pleasures to go for. Jack intrinsically favors all pleasure, whereas Jill withholds intrinsic favor from guilty pleasures. The two disagree not in belief about some special property; but instead, they disagree in attitude.

Now of course a lot happened philosophically in Stevenson's three decades at Michigan, and the influence of his budget line wasn't going to reproduce in me exactly Stevenson's original theory. I'm taken also with A.J. Ayer's original way of putting things - and eventually, Ayer and Stevenson each took on board aspects of the other's approach. My own view isn't either of theirs; it isn't a form of emotivism. But devices that Ayer, Stevenson, and others invented in the mid-1930's turn out, I claim, to be more powerful than even they realized. Such are the joys of hindsight.

The issues now seem much broader than just morality. My teacher and colleague Richard Brandt talked of what it's rational to do or to support. Wilfrid Sellars talked of the "space of reasons". We can broaden the puzzle the great non-naturalists and emotivists addressed, to one of rationality, or to something grandiose like "the place of reasons in a natural world". Reasons are what weigh toward something's being rational. What's puzzling in moral disputes, then, may boil down to what's puzzling about reasons - reasons to do things, reasons to believe things, and the like. What reasons do you have to help others? Do you have reason to care if they suffer, apart from how their suffering comes to affect you? How do you have reason to feel about someone who preys on others? Is the pleasure something brings you always reason to favor it?

Reasons are puzzling, and one thing that's puzzling about them is this: We are living beings, and as such, we are parts of the world of nature. But in the natural world, clearly we're exceptional. Our species has developed refined and ingenious ways of studying the natural world, and these methods tell us a lot about how we are exceptional. For one thing, of course, we are living, and life is so unlike anything else in the universe that it long seemed that the only possible explanation was a special vital principle. Since Darwin, though, we begin to see how aeons of natural selection can account for why life, viewed as part of nature, is so different from non-life. Even among living organisms we're exceptional, and the human brain is vastly more complex than anything else we know about. Human history, politics, social life, learning, and the arts are far more complex than anything even in the life of chimpanzees, though the genetic equipment that allows a human child to grow up to participate in all this is just a last minute evolutionary tinkering, over the past couple of hundred thousand generations, with a tiny proportion of chimp DNA. Biological thinking may give some hints as to how natural selection worked to shape the potentialities of a human infant. We can get some ideal of how babies equipped with these potentialities grow up, in interaction with older people who all started out as babies, to become the human adults we know. Lore, literature, and common sense tell us a great deal, and psychology and social sciences at their best can extend this knowledge and help us integrate it into what we know of the workings of the natural world.

We're exceptional, though, in ways that seem to resist incorporation into any such scientific picture. We have thoughts and opinions and we make assertions to each other. We're conscious of colors and feelings. And we have reasons to do things. Imagine a science of humanity so successful that it could explain, in terms of levels of complexity built on fundamental physics, the sound waves that come from my mouth, all the neuronal patterns in your heads as a result, and all the movements of our limbs and fingers for the next week. Such a science would have to show us as exceptional indeed in the universe. But throughout the era of modern natural science, at least since Galileo and Descartes and Hobbes, crucial parts of philosophy have tackled what's exceptional about us and seems to be left out of the picture. Philosophy is always dealing with how to make sense of new findings in science, thinking how they might transform our visions of ourselves and our surroundings, or how they might fit in with things we always thought we knew. That's by no means all that philosophy does, but questions of what to make of the scientific image form a significant part of our job. And so we ask where in a naturalistic picture of ourselves are beliefs, consciousness, and reasons.

Moore thought that moral facts somehow lie outside the world that empirical science can study. We can broaden this to a claim about the space of reasons as a whole, which, we can say, lies

outside the space of causes. The "space of reasons" is the whole realm of normativity, to use a less picturesque, more technical term that we philosophers like. It's the realm of oughts, we might say, for what I ought to do is what the reasons that pertain weigh toward all told. The reasons to do something, as TM Scanlon puts it, are considerations that count in favor of doing it. Now Stevenson and Ayer devised a cluster of strategies which I want to broaden. (I should warn you that I'll play rather free with their doctrines and motivations, as I've already been doing.) Moral claims, they agreed with Moore, aren't claims that can form part of the empirical sciences. But still, we can understand what we're doing when we make moral claims: according to Ayer, we're expressing emotions or attitudes. There's a broad strategy at work here, a strategy that has come to be known as expressivism. (There's no very good name for this strategy that I know of, but this one, I think, is the least unsatisfactory.) To explain the meaning of a term, to explain the concept that the term conveys, don't offer a straight definition. For normative terms, Moore and the non-naturalists are right that no definition in non-normative terms will capture the meaning. Instead, explain the states of mind that the uses of the term express and don't just explain it as the "belief" that so-and-so. Trivially, normative statements express normative beliefs or judgments. Ayer and Stevenson proposed that moral judgments are feelings or attitudes. I've said that we can broaden the question of meanings to cover normative terms in general. Suppose, then, we try the expressivist twist on oughts in general. What kind of state of mind do [ought] claims express?

Ayer stressed the difference between expressing an attitude and saying that one has it. It is the difference between saying "Boo for lying" and saying "I'm against it." This difference is subtle, since either one of these speech acts gets the hearer to think the speaker is against lying. Stevenson's talk of disagreement, though, lets us get at the difference. If you say "I'm against it", then literally, I disagree if I think you're not against it. If you say "Boo for lying", I disagree only if I disagree with your opposition to lying. What's at issue in the two cases, then, is different: With "I'm against it", what's literally at issue is your state of mind, whereas if you say "Boo!", what's at issue are feelings.

What's at issue with oughts in general, then? we can ask. Jack and Jill need water, imagine, but the hill is slippery. I say that Jack ought now to go up the hill, but you disagree. What's at issue between us? Isn't the issue what to do? Not what to do in your case or in mine, but somehow in Jack's. It's a question of what to do if in Jack's shoes. We explore together what Jack ought to do, engaging in a kind of hypothetical contingency planning. We put our heads together and think the problem through as if on Jack's behalf. Jack himself thinks fleetingly what to do, and decides to follow Jill up the hill. When he falls and breaks his crown and the dangers become more vivid to him, he may come to disagree with that earlier decision. You and I address the same problem as Jack himself rethinks: what to do in his original situation. You disagree with Jack's decision to go up the hill, whereas I agree with it. Switching to normative language, we can describe our states of mind like this: you and Jack both think he ought not to have gone up the hill, whereas I disagree: given the need for water, I conclude, getting the water was, in prospect, worth the danger.

Why should you or I plan, though, for such a fantastic contingency? Why, for that matter, should Jack rethink his decision, when the moving finger has writ and he can't unbreak his crown? It's clear enough why to plan for some contingencies you might face - a traffic jam on

the Dan Ryan Expressway, say, when you want to drive south. Why, though, plan for contingencies you know you won't face - such as Jack's choice of whether to go up the hill? Well of course, mostly we don't. Even if I'm right that ought thoughts are plans, we don't usually worry ourselves with whether, in light of the needs and the Dangers, Jack ought to have gone up the hill. But we do do lots of planning for how to cope with needs and dangers; that's a crucial part of life to plan for. Jack reconsiders after the fact because he'll face such choices again; he's engaged in a kind of rehearsal for further such choices. You and I might join him in this, considering Jack's plight as an exercise in planning for life. Just as Jack might disagree with his earlier decision and so emerge wiser from the calamity, so might you or I. Of course mostly, when we engage scenarios and the places one might hypothetically occupy in them, we aren't thinking to some aforethought purpose - any more than children play to develop their skills and social knowledge, or you read a novel to sharpen your powers of social apprehension. We're just built to engage in such activities, and it's a good thing we are, since doing so functions as rehearsal for later eventualities. We're curious about oughts as well as is.

There's a place in our lives, then, for planning even for the wildest of contingencies. Still, does this really vindicate disagreement in plan? Why treat your plans and mine for Jack's plight as anything on which we could disagree with each other? You have your plans and I have mine; why isn't that just a difference between us as with age, height, or tastes? You have a flatter head than I do, suppose: That's not a disagreement between us; it's just a difference in how we are. You plan, for the contingency of Jack's plight, to stay safe and waterless at the foot of the hill, whereas I plan, as did Jack, to go up the hill with Jill. Isn't this just another difference in our biographies? How is it a disagreement?

Not that there are separate questions of what you are to do in Jack's plight and what I myself am to do: Jack's exact circumstances include everything about him, and our question is what to do if one is he and thus exactly like him in every respect in which we differ. Still, why treat that as something you and I can discuss and agree on or disagree on? Well, I say, that's because we need to be able to put our heads together. Often we need to think cooperatively, treating each other's thoughts like thoughts that occur to oneself, to be considered and supported or refuted, to be accepted or rejected. It is not always good to think alone.

So let's extend Stevenson to say that there is such a thing as disagreement in plan. You and I can disagree on what to do if in Jack's situation, with all his characteristics. This isn't the same thing as "disagreement in attitude" as Stevenson's used the phrase. His disagreement in attitude is disagreement as to what shall happen. Imagine a pacifist who is meek on principle and a bully who takes advantage of this and slaps the pacifist's cheek. When it comes to what shall happen, the two might both favor the same thing: that the pacifist turn the other cheek. On this one point, the two agree in Stevensonian attitude: they both favor the same thing's happening. But they disagree in plan; they disagree on what to do if in the shoes of the pacifist. The bully plans to strike back in such cases, but the pacifist instead turns the other cheek. The bully is planning, to be sure, for the contingency of being someone who is meek on principle, someone who is going, as it happens, to turn the other cheek. But the bully disagrees with the pacifist's plan. The bully's hypothetical preference for the pacifist's situation is to snap out of his pacifism and strike back. The two disagree in plan, then, for the contingency of being the stricken pacifist.

Now the possibility of such disagreement in plan, I claim, has far-reaching consequences. I'll sketch a few of them, though I won't really be able to argue for what I say; I want rather to contemplate what happens if these claims are true. The first chief consequence is that we can deal with complex normative claims. We get an answer to the Frege-Geach challenge to expressivism. The mother admonishes, "If lying is wrong, then getting your little brother to lie is wrong too." Starting with the notion of disagreement, we can say canonically what the content of such a plan-laden claim is. The mother has come out in disagreement with any plan to shun lying but get little brother to lie. In general, to get the content of a plan-laden claim, we map all the combinations of pure plans and pure factual beliefs with which the claim is in disagreement. Disagreement is the key to content; content is what there is to agree or disagree with. So allowing for disagreement in plan gives us plan-laden content - normative content.

The second chief consequence will sound surprising, coming from an expressivist. In a sense, it follows, normative terms like 'ought' refer to properties and relations - indeed to properties and relations that are natural, that can figure in an empirical science of humanity. My argument for this is transcendental: As planners, capable of agreement or disagreement in plan, we are each committed to this naturalistic-sounding thesis. Once we establish this thesis as one to which we are all committed, this thesis of natural constitution, we can proceed to assert it: There is a natural property that constitutes being what one ought to do. Thus we're all committed to agreeing, in a way, with normative naturalists: the term 'ought' refers, in a sense, to a natural property.

And what property is this? That's not a linguistic question; it's the grand, basic question in ethics, the question of how to live. You accept an answer to this question if you have fully thought out what to live for and come to a conclusion. Consider a view that fits some aspects of Henry Sidgwick's doctrines: A universal hedonist whom I'll call Henry plans always, in every conceivable contingency, to do whatever holds out maximal prospects for net pleasure in the universe. Henry, then, has a view about the property that constitutes being what one ought to do. It is, he says, the property of being unihedonic, as we might call it: the property of holding out maximal prospects for net pleasure in the universe.

Henry, then, accepts this thesis of natural constitution. Indeed not only does he think that there's a natural property that constitutes being what one ought to do; he has a view on what it is. Many of us, though, don't have anything like a complete contingency plan for what to pursue in life, or a formula for constructing such a plan. Still, I claim, we are each committed to the thesis of natural constitution. For suppose you are at least consistent. Then the thesis is something you'd accept if, fantastically, you completely filled out your views on how to live, and did so without changing your mind about anything. Any way of filling out your plans, becoming hyperdecided on how to live, brings with it accepting the thesis of natural constitution. So it's something you are already committed to as you think your way toward a fuller view of how to live. It's something that obtains, you can say, no matter what turns out to be the way to live.

Now I don't mean you to be convinced by this cryptic sketch of an argument. Even if I had succeeded in making the argument clear, it would raise many issues I can't quickly resolve. I want to sketch the possibility, though, of a view of normative concepts that has us sounding like expressivists, like non-naturalists, and like naturalistic realists in important respects - all at the

same time. We start out with devices of the classic emotivists: with disagreement in plan reminiscent of Stevenson, and with Ayer's talk of expressing a state of mind. We let the state of mind in question be a kind of contingency planning for living. As Ayer and Stevenson saw, we derive Moore's conclusion that normative concepts aren't naturalistic. Two people might agree, in naturalistic terms, on all the natural facts and still disagree basically in plan. There's something at issue between them, but not something we can put in naturalistic terms. It's a question of how to live. Still, as naturalistic realists insist, normative terms like 'ought' do signify natural properties. That's something that Ayer and Stevenson didn't say, but it falls out as a consequence of some of their ways of thinking.

Simon Blackburn coined the term 'quasi-realism' for a program like this one. We start out without helping ourselves to ethical and other normative properties. But then we earn the right to speak as realists do. Indeed we may be hard pressed to identify any real differences between naturalistic realism, non-naturalistic realism, and expressivism, once these positions are suitably refined. We may have a happy convergence of different approaches to metanormative theory. I think of what I'm sketching as filling out this program that Blackburn proposed. My impression is that he's skeptical of the extremely metaphysical-sounding claims that I've been sketching here, but if those claims are right, then perhaps they fit Blackburn's program.

Everything I've been saying depends on a distinction that's been in the air in recent decades but which wasn't much around when Ayer and Stevenson were doing most of their work. It's the distinction between properties and concepts. The property of being water, we can say, turns out to be the property of being H<sub>2</sub>O, of consisting in molecules of a certain kind. Still, the concepts are different: the prescientific concept of being water isn't the scientific concept of being H<sub>2</sub>O. It was a live question at one time whether water was H<sub>2</sub>O, a question on which people could coherently disagree. People disagreed as to whether water is H<sub>2</sub>O; they didn't disagree as to whether water is water. We can ask what was at issue. Disagreement, then, is a matter of concepts, not properties: it isn't always preserved when we substitute distinct concepts of the same property.

Once we have this distinction, we can say this: All properties are natural, but some concepts of properties aren't descriptive and naturalistic. Some concepts find their place not in naturalistic description but in planning. Suppose, then, that Henry the universal hedonist is right on how to live: the thing to promote in life is the happiness of all. Then the property of being what one ought to do just is the property of being unihedonic, of holding out maximal prospects for total net pleasure in the universe. But the concept of ought is distinct from the concept of being unihedonic. For a perfectionist Percella can dispute with Henry: Perce says that the unihedonic thing isn't always what one ought to do. Henry understands her - and nothing about logic or our linguistic conventions by itself settles who is right. Perce and Henry have the same concepts; that's why they can engage each other's claims and not just talk past each other. Henry is right, we're supposing, and so the terms 'ought' and 'unihedonic' refer to the same natural property. But conceptually, Perce is coherent. Once she explains what perfection consists in, on her view, we know what's at issue between her and Henry. It's whether to live for universal happiness or to attain that kind of perfection.

This scheme, as I've been saying, has attractive features. Some tenets of Sidgwick's and Moore's ethical intuitionism seem hard to escape, and the scheme delivers these tenets. The inescapable tenets consist, it turns out, just in what we'd have to accept if we are to plan our lives coherently and intelligibly. We don't need non-natural properties, just the kinds of non-descriptive, non-naturalistic concepts that would have to figure in planning. Normative concepts do signify natural properties, we can say, but they have their own special way of doing so. The scheme respects normative thinking: it avoids any blanket debunking of it - though we should still debunk certain theories of what normative thinking consists in. And it's a good thing that we can see normative thinking as inescapable in intelligent living. For normative thinking figures in a wide range of areas that we couldn't give up as nonsense. Normative epistemology, for instance, we can now say, consists in contingency plans for forming beliefs. It's a serious question, for instance, whether the evidence supports a Darwinian theory of natural selection, and more broadly, what the canons of scientific evidence are. I take these to be planning questions, questions of how much credence to put in theories given various epistemic contingencies. They are questions of what we ought to believe. Oughts are to be found even in places far from ethics.

Is this picture I have given, though, a naturalistic one? Does it really let us dispense with all mumbo jumbo of a non-natural realm? Not exactly. We can view ourselves as complex products of natural selection and the kind of cultural history that natural selection could make possible. We can see, in these terms, why beings like us might be interpretable as planners who share our planning thoughts. Suppose we view ourselves this way, and suppose furthermore, we interpret such natural beings as keeping track of what disagrees with what. Then we are interpreting ourselves as having normative thoughts. We can see, in short, why natural beings like us would be plausibly interpretable as having normative thoughts.

The scheme I've been sketching has a further happy consequence: If you start out as a non-naturalist, you have to accept certain features of the space of reasons as just brute normative facts: for instance, that the normative supervenes on the natural. Once we see normative facts as plans, we see why this supervenience is something that any planner is committed to. Plans must be couched in empirical, naturalistic terms because we have to be able to recognize the situations the plans address. A plan to do whatever there's most reason to do, for instance, is no plan at all, until it's supplemented by an account of how to recognize what there's most reason to do. With this supplement, the plan is in effect couched in empirical terms.

On the other hand, the scheme in no way lets us substitute naturalistic thinking for normative thinking. Instead it follows Moore in concluding that there's just no substitute for normative thinking. And moreover, it doesn't tell us how to translate, in strictly naturalistic terms, claims about people's normative states of mind. Take the claim, "Jack is convinced that he ought to go up the hill." I haven't indicated how to translate such a psychological claim into terms that fit a broadly Galilean picture of the universe. Imagine we understood Jack completely as a physical system. Imagine we understood him, at many different levels of explanation, as a product of natural selection and a vastly complex human ecology. This would include grasping the explanatory patterns of his neurophysiology, understanding how evolutionary signaling theory applies to his patterns of neural firings and the sound waves that come out of his mouth, and all sorts of things like that. My hope is that the expressivistic scheme I've sketched would then let

us see why Jack, so viewed, is conveniently interpretable as thinking that he ought to go up the hill.

He'd be conveniently interpretable that way, I'm saying. For Jack, viewed as a natural system, is conveniently interpretable as keeping track of his surroundings. (We have some naturalistic idea, for instance, how rats keep track of their position in a maze.) He'll be conveniently interpretable as planning, and we can conveniently interpret him as agreeing and disagreeing with combinations of plan and mundane fact. And that, I'm saying, is all we need if we're to interpret him as having normative thoughts.

I'm speaking of convenient interpretation, how a natural being, viewed naturalistically, might be conveniently interpreted. But how much does this establish if true? For anything I've claimed, a convenient interpretation might be no more than a convenient fiction - like the stupidities we attribute to the computers on our desks. When Jack is conveniently interpretable as thinking he ought to go up the hill, is that what he's really convinced of? Is he really thinking he ought to go up the hill?

That, I've implied, is a question of agreement and disagreement: for issues of meaning and interpretation, as I've been harping, agreement and disagreement are the key. What Jack accepts by way of oughts and iss, I've been saying, is a question of which possible states of mind he disagrees with and which he doesn't. Now I haven't offered any naturalistic translation of claims about disagreement. And if I did, my translation might be subject to Moore-like challenges. Do Jack and Jill really disagree with each other on Whether they ought to go up the hill? Suppose you and I disagree on this question of how to interpret them, this question about disagreement. What's then at issue between us? That's Moore's challenge, transferred from ethics to the theory of meaning itself. It's a question about the meaning of meaning, or the meaning of claims about mental content, claims about what people are thinking. And I haven't so much as sketched how to respond.

Issues of meaning and mental content may in part themselves be normative issues. A number of leading philosophers have asserted that they are, and whether meaning is in some sense "normative" is a daunting question. The question has received intensive scrutiny over the past decade or two, and the issues still aren't entirely clarified. Perhaps to understand claims about what Jack really is thinking, you have to understand about commitments, or about correctness - and the concepts of commitment and correctness seem to be normative ones. Take two claims that contradict each other: Jack thinks, imagine, that snow is white whereas Jill thinks that nothing is white. Jack and Jill disagree - and this implies, among other things, that we ought not to accept both these claims at once. We ought not both to accept that snow is white and to accept that nothing is white - that's a normative claim. Claims about states of mind and their content seem themselves, then, to be fraught with ought. Some philosophers argue that this appearance dissolves on close scrutiny, and I'm not claiming to establish this "normativity of meaning" thesis or even urging us to accept it. But I'm not denying it either. Perhaps the right theory of normative thinking must itself be a normative theory. Perhaps Robert Brandom is right that it's norms all the way down.



If so, then the account we end up with won't be strictly expressivistic or strictly quasi-realistic, by a stringent standard of what qualifies under these terms. That is to say, it won't fit the following pattern: that it starts out helping itself to a purely non-normative reality, and ends up, all on its own, earning our right to realistic ways of talking about oughts.

Still, the account is expressivistic in a weaker sense. It draws on central philosophical devices of Ayer and Stevenson. And if it succeeds in its ambition, it makes clear how natural beings like us would be conveniently interpretable as having ought thoughts. As for whether this interpretation would really get things right, perhaps we should take this question with a grain of salt. Suppose none of our uncertainties were scientific: we understood Jack completely in purely naturalistic terms, insofar as beings like him or us can be understood in naturalistic terms. Suppose you and I none the less have competing interpretations of Jack, and that these interpretations both are as convenient as can be. What's then at issue between us? as Moore might ask. Well, perhaps nothing real: perhaps nothing real is at issue. That's a familiar enough conclusion in the past half century of philosophy, with Quine and in his wake. We can perhaps be skeptics about picky questions of meaning that go beyond questions of convenient interpretability.

With normative questions, in contrast, it's hard to be a skeptic; it's hard to take the questions with too many grains of salt. The question of what to do is inescapable. Sartre's man who asks himself whether to join the resistance or take care of his mother can't dismiss the question as nonsense. And when he comes to a decision, he has accepted an answer to a normative question, the question of what he ought to do in the circumstance.

Or at least, he has come to a normative view if a big if is satisfied. He has reached an ought conclusion if there is such a thing as agreement and disagreement in plan, if we can come to agree or disagree with his conclusion. Whether there is such a thing as disagreement in plan, though, is a deep question - as I have indicated. You and I can certainly think what to do if in this Frenchman's exact circumstances, and form a different plan from his. The deep question is why to treat this as any sort of disagreement. Why think there is such a thing as disagreement in plan? My answer has gone in two stages. First, in planning I have to be able to change my mind, and this amounts to disagreeing with things I had concluded earlier. Second, in thinking how to live, we need each other's help.

Disagreement in plan, I've been saying, is the key to explaining normative concepts - that along with Ayer's distinction between expressing a state of mind and saying that one is in it. The concepts we explain with these devices act much as the classic non-naturalists recount. Explanatory devices we get from the classic emotivists Ayer and Stevenson, then, lead us to crucial aspects of Sidgwick, Moore, Ross, and Ewing. Now if we really get this much convergence, that should be grounds for celebration: perhaps we're really getting at what's going on with is and ought. Of course, it is bound to go on being controversial whether we do achieve this convergence - and legitimately so, as we work to understand better the tangle of issues in play. The convergence also leaves the question, though, whether an expressivism that draws on Ayer and Stevenson tells non-naturalistic normative realists anything they didn't know before. So let me review some of the ways that Ayer's and Stevenson's devices lead to illumination, if I'm right. (I've been playing on my institutional tie to Stevenson, whereas I don't know if Ayer ever set foot in Ann Arbor. But Stevenson, it's said, was fired from Yale for the immorality of

his theories, and the first time I ever set eyes on Ayer, across a room, his doctor, apparently, had forbidden him to travel to Yale to lecture. "I know I've got to die some day," he declared, "but not in New Haven!" I don't know if Stevenson would have welcomed the sentiment if he had still been alive.) Anyway, first of all, Ayer's and Stevenson's devices let us take what comes across a mystery, as Moore presents it, and see it in terms of something familiar and pretty comprehensible. We can explain ought convictions as plans, oughts as deliverances of planning. Second, the brute features of Moore's non-natural realm fall out as things a coherent planner would have to believe in. We get supervenience of the normative on the natural, and get something that fits Moore's talk of the good" as something natural: there's a natural property, I've been saying, that constitutes being what one ought to do.

So do we eliminate the mumbo jumbo of a non-natural realm we can intuit? Not exactly, but we see why a being like you or me would have to be interpretable as committed to this mumbo jumbo. We work toward a naturalistic view of why we'd have to be so interpretable. And all this is in a world where all properties are natural - though non-naturalistic concepts apply to it, and we can see why.

We ourselves are parts of the natural world we study, and the moral, perhaps, is that this makes for concepts that aren't just naturalistic classifications of nature. I've been exploring some ways all this might happen - but mostly, I've been musing over the consequences of a philosophical approach, a theory of normative concepts. I've gone in haste over a number of theses and issues. I've sketched this approach not so much with an eye to laying out "What does it mean and how do we know?" but with an eye to the question "So what?" In particular, what does all this say about whether we live as purely natural parts of a purely natural world? So is what I've been sketching, you might ask, naturalism, non-naturalism, or something else? It's a view that's all three, I answer, in ways we need to distinguish. It's not the classic emotivism of Ayer and Stevenson, but still, devices those thinkers invented help us construct a view that takes in crucial aspects of all three.