

The John Locke Lectures 2009

Being Realistic about Reasons

T. M. Scanlon

Lecture 2: Metaphysical Objections

As I said in my first lecture, the idea that there are irreducibly normative truths about reasons for action, which we can discover by thinking carefully about reasons in the usual way, has been thought to be subject to three kinds of objections: metaphysical, epistemological, and motivational or, as I would prefer to say, practical. Metaphysical objections claim that a belief in irreducibly normative truths would commit us to facts or entities that would be metaphysically odd—incompatible, it is sometimes said, with a scientific view of the world. Epistemological objections maintain that if there were such truths we would have no way of discovering them. Practical objections maintain that if conclusions about what we have reason to do were simply beliefs in a kind of fact, they could not have the practical significance that reasons are commonly supposed to have. This is often put by saying that beliefs alone cannot motivate an agent to act. I think it is better put as the claim that beliefs cannot explain action, or make acting rational or irrational in the way that accepting conclusions about reasons is normally thought to do. I will concentrate in this lecture on metaphysical objections.

I do not believe that irreducibly normative truths have untenable ontological implications. To explain why I think this, it will be helpful to begin with some general remarks about ontology, starting with Quine. In his famous essay “On What There Is”

Quine proposed that we understand what he called our “ontological commitments” in the following way. The ontological commitments of a set of statements are determined by first translating these statements into the language of first-order logic, and then determining what things there must be in the “universe of discourse” of a model in which all of these statements are true. These things are what we are ontologically committed to in accepting those statements. What Quine later called “ontological relativity” arises from the fact that there will be more than one way of translating any set of statements into “regimented form,” and these different translations may yield different ontological commitments.

So far, Quine’s view says nothing about what our ontology should be. It is compatible with what might be called the permissive first-order view, according to which we should decide what sentences to accept by applying the criteria appropriate to the relevant first-order disciplines—empirical science, mathematics, and so on perhaps including (as I would say but Quine probably would not) our best moral thinking—and our ontology is simply the set of ontological commitments that these sentences have, determined by Quine’s method. But that method is equally compatible with various more restrictive views, according to which we should, for example, avoid ontological commitment to anything other than physical objects, or should limit our ontology as much as possible, and should reject statements that would have ontological commitments that violate these strictures, even if they are attractive on first-order grounds. In his early years, Quine seems to have held some restrictive views of this kind. In his paper with Goodman, they reject ontological commitment to abstract entities, and, more generally, he seems to have favored a kind of ontological minimalism (“a taste for desert

landscapes.”) But by the middle 1960’s, when he wrote *Set Theory and its Logic*, these views had moderated. He still shows some preference for limiting one’s ontology, other things being equal, but he is willing to trade off this concern against such concerns as the ease and shortness of proofs. The willingness to accept such trade-offs does not seem to me to be consistent with any very serious doctrine of ontological minimalism or the rejection of abstract objects as ontologically unacceptable.

Whatever Quine’s view may have been, I myself hold and will try to make plausible to you a permissive first-order view. We should decide which statements to accept on the basis of the best first-order reasoning appropriate to them, that is to say, forms of reasoning such as scientific reasoning, moral reasoning, and practical reasoning about reasons more generally. These first-order domains are not always autonomous: if the claims of one domain can conflict with those of another, then these claims need to be reconciled, and some of them perhaps given up. For example, we might have a first-order theory of witches and spirits. That is, we might have established criteria for deciding whether someone is or is not a witch, and whether or not a ghost is present. But the conclusions of these theories would entail claims about events in the physical world and their causes. These claims would therefore conflict with those of physics and other empirical sciences, and as a result of these conflicts these theories—and the idea that there are witches and ghosts—should be rejected. But this rejection is based entirely on first order grounds, in this case scientific grounds.

What this “permissive first-order” view excludes are general ontological constraints, such as a general preference for ontological economy, or a general stricture against commitment to abstract objects. There may be good reasons for preferring

simpler or more economical theories to more complicated ones. For example, it may be good scientific practice to prefer simpler physical theories. But when there are such reasons, they arise, as in the case just mentioned, within particular first order domains. The rationales for them, and what counts as “simplicity” in the relevant sense, will be specific to those domains.

It is sometimes said that the thesis that there are moral facts, or irreducibly normative truths, is incompatible with a scientific view of the world, but I see no reason to believe that this is the case. Science is a way of understanding the physical world—the world of space and time, physical objects and causal relations. Belief in irreducibly normative truths would be inconsistent with science if—like beliefs in witches and spirits—it committed one to claims about physical objects and causal relations that science denies. But belief in normative truths has no such implications. It may seem that the role of normative beliefs in motivating action, and the possibility of our coming to know normative truths involve claims about causal interaction. I believe that this is not so in either case, as I will explain later lectures dealing with the epistemological and practical objections.

It is important here to distinguish between *the universe* and the logical notion of a “universe of discourse.” Science is a way of understanding the universe—the natural world. Its conclusions represent our best understanding of what that world contains and what happens in it. A universe of discourse is merely formal: a way of representing all those things that are presupposed by some set of statements, about the natural world or anything else. Accepting science as the way of understanding the natural world entails rejecting claims about this world that are incompatible with science, such as claims about

witches and spirits. But accepting science does not mean that this world and the things in it are the only things we can refer to and talk about sensibly. It does not entail any limitation on our universe of discourse.

This view may seem to entail that some things that must be counted as part of our universe of discourse “exist in a different sense” than other elements in that universe, such as tables, chairs, mountains and electrons. This may seem to run contrary to the doctrine, attributed to Quine, that ‘existence’ is univocal. According to the first-order view I am proposing ‘existence’ is univocal in that what it expresses is captured by the existential quantifier and the logical rules governing its use. But what is required to justify any existential claim, and what follows from such a claim, varies, depending on the kind of thing that is claimed to exist. The claim that mountains exist is licensed by and licenses certain other claims about the physical world. The claim that there exists a number or set of a certain kind is licensed by and licenses certain other mathematical claims. The claim that a right exists is licensed by and licenses certain other moral claims. And in each case that is all there is to it.

One objection to the view I am taking might be that it is too permissive. According to this view, it might be said, we could adopt some way of talking which allowed for existential generalization and, as long as this way of talking was well-defined, internally coherent, and *did not have any presuppositions or implications that might conflict with those of other domains, such as science*, we would be committed to the view that the things quantified over in this way of talking exist. They would be among our “ontological commitments”—that is to say, part of our universe of discourse. But the fact that these things are insulated from other domains mentioned in the italicized

phrase may just ensure that talk of the items in question is perfectly pointless. Can we take seriously an idea of existence that comes so cheaply?

My answer, first, is that the question to be asked about the entities in question is not whether they really exist, but whether we have any reason to talk about them. If we do, then there is reason to count them among our ontological commitments and part of our (purely formal) universe of discourse, and (given the assumption of independence that I mentioned) no reason not to. There is (given this assumption) no further question of existence.

It might be said that on this view “ontological commitments” are not *ontological* in a serious sense. Perhaps this is correct (and a consequence of accepting the idea that the univocal meaning of ‘exists’ is captured entirely by the logic of the existential quantifier.) What more serious sense of ontology is there? In my view the only relevant criteria of existence arise from the particular first-order domains that we have reason to take seriously: existence in the natural world described by science, mathematical existence, and so on.

The position I am taking here is similar to the one famously (notoriously?) advocated years ago by Carnap, and famously criticized by Quine.¹ It may help to clarify what I am saying if I explain how my account is like Carnap’s and how it differs from his. Carnap distinguished between questions of existence that were “internal” to a framework,” such as mathematics or physics, and questions that were “external” to such

¹ Carnap, “Empiricism, Semantics, and Ontology,” *Revue Internationale de Philosophie* 4 (1950), reprinted (revised) in Benacerraf and Putnam, pp. 241-257. Quine, “Carnap’s Views on Ontology,” *Philosophical Studies* 2 (1951), reprinted in Quine, *The Ways of Paradox and Other Essays* (New York: Random House, 1966), pp. 126-134. Page references to both articles will be to the reprinted versions.

frameworks. For example, the question of whether there exist numbers of a certain kind, such as twin primes, or whole numbers that are solutions to a particular equation is an “internal” mathematical question, whereas the question whether numbers really exist is an external question. Carnap held that only internal questions were genuine theoretical questions. External questions, understood as questions of fact about the existence of certain entities were, he said, “pseudo questions.” The only genuine external questions, he thought, were practical questions about whether or not to adopt a particular framework. Quine objected that the distinction Carnap sought to draw between internal and external questions was really an instance of the analytic-synthetic distinction. It is not only “external “ questions that involve the choice of a “convenient conceptual scheme or framework for science.” According to Quine, such choices are at issue in the consideration of any scientific hypothesis.²

Carnap spoke of “frameworks.” I have been speaking of first-order domains or practices, such as mathematics, science, and moral and practical reasoning. The idea of a domain is, admittedly vague. My understanding of the idea is driven simply by the examples I have just given. A domain is determined by a set of concepts that it deals with, such as number, physical object, or morally right action, a certain number of things taken to be settled truths that employ this concept, and accepted procedures for settling questions employing these concepts. What these procedures are is often vague or incomplete. In contrast to Carnap, I do not take these procedures to be determined by “linguistic rules” for the use of the terms in question. People can all use terms such as

² Quine, 134. For a fuller analysis of the debate between Carnap and Quine than I can under take here, see Stephen Yablo, “Does Ontology Rest on a Mistake?” *Proceedings of the Aristotelian Society*, supp. vol. LXXII (1998), pp. 229-261.

‘number,’ ‘set,’ or ‘wrongness’ without misuse of language even though they have some disagreements about the best ways of determining the facts about such things.

In discussing the possibility of an isolated domain, one that had no implications for other domains, I said that the only external question about the entities it referred to would not be “whether they really exist, but whether we have any reason to talk about them.” This sounds very much like Carnap, and I do agree with him that *in some cases* the only external questions are practical ones. But this is not always so. Meaningful external questions about a domain, in my view, are questions about whether the claims and presuppositions of that domain are in fact fulfilled. Most commonly they are questions about whether procedures internal to that domain can actually deliver conclusions that have the significance for us that they claim or presuppose. So, for example, the conclusion that someone is a witch, or that there is a spirit in the room, have the significance they claim only if they have implications for what happens in the physical world. So there is an external question about whether the methods that are claimed to establish such conclusions actually support these implications. (I do not think Carnap would disagree.) This is a genuine question, answered, I believe, in the negative. It is not a metaphysical question, however, but a question of physics.

This example indicates one source of unclarity in the internal/external distinction. The first-order claim that someone is a witch not only presupposes, in order to have the significance that such claims are supposed to have, that witches have certain causal powers. Claims about what witches cause or have caused are, presumably part of the first-order theory of witches. So it would seem that such causal claims are internal claims. What I have been treating as external is the question whether the methods of reasoning

that, I am supposing, the first order theory of witches takes to be sufficient to establish that someone is a witch actually support these causal claims. I call these questions external because they concern the adequacy of these methods.

Physics itself, and our everyday understanding of physical reality, may employ notions that require further explanation that may count as external. For example, there are questions about how the ideas of cause and of a law of nature are to be understood. Insofar as these are questions that cannot be answered *by* science, because they concern the content and significance of scientific conclusions themselves, they would count as external in the sense I have in mind: they arise from a first order inquiry but cannot be answered satisfactorily within it. Perhaps there are similar questions about whether physical objects endure or perdure. I am not denying that such questions make sense, or claiming that they are shallow, or that they are not theoretical questions but rather pragmatic ones. So I am not saying any of these things about all claims that would be normally counted as metaphysical or ontological.

Morality, too, has presuppositions. Talk about moral right and wrong presupposes that there are moral standards that everyone has good reason to take seriously as guides to conduct and as standards for objecting to what others do. But the ordinary ways of understanding morality do not make clear what these reasons are, and there is therefore an external question whether there are such reasons—whether the usual ways of establishing that a form of conduct is wrong also guarantee that there are good reasons not to engage in it. This question is not metaphysical, or scientific, but normative, hence “practical” in one sense of the term—a question about what we have reason to do. (But it

may not be a practical question in the sense that Carnap had in mind, if what he meant was a question about whether it would be *useful* to adopt a particular framework.)

Although not all external questions are, in my view, practical questions, there is a sense in which all such questions are driven by concerns that might be called practical in a broad sense. They are questions about what is required in order for the methods we commonly use for reaching conclusions in a given domain to have the significance they claim, or that we commonly give them. They thus depend on our concerns, in a broad sense.

One objection to the view I have been advancing is that the priority it gives to what I have called “first-order” domains is arbitrary and unmotivated. I have specified these domains simply by a list of examples—natural sciences, mathematics, moral and practical reasoning. I called them first-order to distinguish the modes of thinking they involve from “higher-order” questions about them—such as about the adequacy of the modes of thinking they involve to deliver conclusions with the significance that is claimed for them. But what makes a set of questions and way of resolving them “first-order”? I have rejected the Carnapian idea that these ways of thinking are enshrined in linguistic rules, and it would be clearly inadequate to say just that these are the modes of thought that we customarily employ.

So, why just these familiar domains and not others? In particular, why not a “first-order domain” of metaphysics, or ontology, which is concerned with the scope and nature of our overall ontological commitments? If, as I have said, a first-order domain, such as witch theory, may be rejected because its implications or presuppositions conflict with the first-order domain of natural science, why shouldn't other domains such as morality

be open to possible limitation, or even rejection, when they conflict with (our best) metaphysics?

My answer is the one suggested by what I said about the possibility of an isolated domain with ontological commitments but no implications or presuppositions that conflicted with those of other first-order domains. The question about such a domain, I said, was whether we have reason to be concerned with the conclusions it delivers. Similarly, the question about an autonomous first-order domain of ontology is whether we have reason to be concerned with the questions it addresses and the answers it yields. My view is that our universe of discourse is a purely formal notion, not “a world,” let alone “the world,” and that our only reasons for being concerned with what this formal universe contains are ones that arise from the particular domains that contribute to it. We have no domain-independent reason to be concerned with how many things are quantified over in all of our first-order domains taken together, or with whether these things are abstract or concrete.

The implications of this for the idea of a first-order domain of ontology might be put in either of two ways. One is that there is no such domain—that is, that we have no reason to be concerned with the questions it would address. The other is that there is such a subject, and I have been engaging in it in this lecture. But my ontological view is that ontological questions are settled by the other particular domains we have reason to be concerned with (and the interactions between them.) That is all there is. I am inclined to think that these two ways of describing the situation come to the same thing. So you can take your pick.

The view I am defending helps to make sense of the place of “indispensability” arguments, such as the idea that we should accept ontological commitment to numbers because they are indispensable for science. Such arguments seem to be given a certain amount of credence. Hartry Field, for example, goes to a great deal of trouble to show that ontological commitment to numbers is not actually needed for science. He would be wasting his time unless indispensability had some serious ontological implications. But the idea that it does have such implications seems odd. (As I have already remarked in commenting on Quine’s later views about sets.) Does existence really depend on what is useful for us? The existence of physical objects does not depend on our interests, or our practical or intellectual needs, and I don’t see how the existence of abstract objects could do so either.

According to the view I am advancing, the existence of numbers and sets is determined by mathematical criteria. Assuming that these criteria are coherent and sufficiently determinate, there is no further question about whether these entities exist. What the indispensability of mathematics for science does is to give us one kind of reason to be concerned with mathematics, and in particular with whether there exist numbers of various kinds.

I said earlier that I did not believe that the problems raised by moral facts, or by irreducibly normative truths, are properly described as ontological. The insignificance of the perfectly general idea of ontological commitment provides one reason for saying this. Another reason is that, contrary to what is sometimes said, belief in irreducibly normative truths does not involve commitment to any special entities. The essential element in normative statements is not a term referring to an entity, but a relation: the relation $R(p,$

c, a), that holds between a proposition, a set of conditions, and an action or attitude when p is a reason for a person in situation c to do or hold a.³

This formulation of the basic normative relation is intended to be non-committal on important normative issues. It is consistent, for example, with the view that the reasons an agent has depend on his or her desires, because it leaves open whether c contains facts about the agent's desires. Also, the fact that the relation R contains no place for the agent him or herself may seem to entail a commitment to the view that all reasons are *general*—that something is a reason for an agent only in virtue of certain facts about his or her situation, and is a reason for that agent only if it is also a reason for any other agent in similar circumstances. This is indeed my view, but it is not entailed by the formulation I have given. That formulation allows for the possibility that the particular identity of the agent may figure among the conditions c, in virtue of which p is a reason for him or her to do a.

The things denoted by terms occupying the first position in a statement, R(p, c, a)—the things that are reasons—are not some special kind of normative entity but ordinary facts, usually facts about the natural world. So, for example, the fact that a piece of metal is sharp, is a reason to use it in order to cut something, and under most conditions a reason not to press one's hand against it (unless other factors give one reason to cut one's hand.) The distinctive aspect of normative truths is thus a matter of what

³ A proposition, p, cannot be a reason unless it is the case that p. So it would be natural to say that the things that are reasons are facts, such as the fact that p. But since, as I will say later, we can discuss whether p *would be* a reason if it were the case that p, the essential normative content of the reason relation seems to be independent of the truth of its first component.

Quine called “ideology” (the predicates we employ) rather than ontology (the things we quantify over.)

This will not comfort many of those who have objected to such truths. John Mackie’s famous metaphysical objection to objective moral values was not just to special entities, but also to moral “qualities or relations” which, he said, would be “of a very strange sort, different from anything else in the universe.”⁴ According to Mackie, the claim that there are “objective values” would involve the claim that certain actions “have to-be-doneness built into them,” or that certain situations would have “a demand for such-and-such an action somehow built into it.” The claim that there are such things, Mackie said, “is not just meaningless, but false.”⁵

In fairness to Mackie, I should emphasize that, like most people discussing these issues at the time he was writing, he was concerned with morality, not with practical reasons more generally. When he speaks of claims about objective values, he may intend to contrast these with claims about “subjective” values—claims about what a person ought to do, or has reason to do, that, unlike moral claims, are claimed to hold only insofar as the agent has certain desires or aims. Mackie may have no objection values, or claims about reasons, of the latter kind.

If so, however, his position suffers a certain instability. As I have pointed out in my first lecture, the claim that a person has reason to do what will promote the satisfaction of his or her desires is itself a normative claim. Indeed, it is an “objective” normative claim, since it does not *itself* depend on what people desire, or on what aims they have. If there is something *metaphysically* odd about objective normative truths,

⁴ *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin Books, 1977), p. 38.

⁵ *Op. cit.*, p. 40

then this supposed truth (that people have reason to do what would satisfy their desires, or promote their aims) is just as odd as any other. The disagreement between someone who thinks that all reasons for action depend on the agent's desires and someone who thinks that there are some reasons that do not depend on agents' desires is a *normative* disagreement, not a metaphysical one. So Mackie's "argument from queerness," insofar as the queerness involved is metaphysical, is an argument against irreducibly normative truths or any kind, not just objective moral values. This is how I am going to take his argument, I hope not unfairly.

When Mackie says that there is nothing "in the universe" like the normative relation I have been describing, what does he mean by "the universe?"⁶ We should note a distinction here, parallel to the distinction mentioned earlier between "the (physical universe)" and "the universe of discourse" of things to which we are ontologically committed. In that case the distinction was between two sets of objects: those that exist in the physical world and those formally included in the universe of discourse. In the present case the distinction is not between two sets of objects but two collections of facts: those that comprise the natural world and those that are simply the reflection of all the things that are quantified over in statements that we accept as true.

⁶As John McDowell observed, Mackie's argument "involves a tendentious use of 'the world.'" "Values and Secondary Qualities," in Ted Honderich, ed., *Morality and Objectivity* (London: Routledge & Kegan Paul, 1985), p. 185, n36. See also the beginning of McDowell's "Non-Cognitivism and Rule Following." Those holding positions more similar to the one I favor do sometimes state their claims in terms that may invite Mackie's response. Crispin Wright, for example writes that if natural number is a sortal concept then "its instances, if it has any, will thus be *objects*, furnishings of the world every bit as objective as mountains, rivers and trees." He goes on to say that the truth of the appropriate arithmetical statements insures that this concept does indeed have instances. *Frege's Conception of Numbers as Objects* (Aberdeen University Press, 1983), p. 13. Although I agree with much of what Wright says, I would not put it in this way.

Normative truths such as the claim I just mentioned about sharp objects are “about the natural world” in one sense: they make claims about a relation among things that are in the natural world: facts (such as the sharpness of the metal), agents in certain situations, and actions these agents might perform.) When Mackie denies that natural facts have “ought to be doneness built into them” he might simply be denying that this relation ever holds—that facts ever provide reasons (denying that, for example, the fact that the metal is sharp is ever a reason, or a reason independent of one’s desires, not to press one’s hand against it.) He may indeed be denying this. But he seems not simply to be denying this first-order normative claim. Rather, he seems to be claiming that such first order claims are false *because* (understood as irreducibly normative claims) they have presuppositions that cannot be endorsed. Specifically, such claims are false because the relation R itself is not a property “in the world” and the fact that it holds of certain facts and actions is not a fact “in the world.”

If by “the world” one means the natural world of physical objects and causal relations, which science aims to describe, then there is no disagreement. Those of us who believe in irreducibly normative truths would not claim that the normative relation R itself is part of the (natural) world—that to claim that it holds is to make a claim about natural facts. Indeed, we explicitly deny this. Normative facts about reasons, as we understand them are “part of the world” only in the broader sense in which “the world” is simply the reflection of all true sentences. Normative claims as we understand them are thus “incompatible with a scientific view of the world” only if, in addition to holding that everything in the natural world can be explained by science, such a “scientific view”

holds that nothing other than this world can be the subject of true statements. As I have said, science does not entail this.

It is perhaps worth noting that some contemporary philosophers who are not realists about the normative accept that normative statements may state truths about the world in this broader sense. Allan Gibbard, for example, says that if by “facts” we mean simply “true thoughts,” then there are normative facts. If there “is no more to claiming ‘It’s true that pain is bad’ than to claim that pain is bad; the fact that pain is bad just consists in pain’s being bad; [and] to believe that pain is bad is just to accept that it is,” “Then it’s true that pain is bad and it’s a fact that pain is bad.”⁷ Similarly, Simon Blackburn writes, “There is no harm in saying that ethical predicates refer to properties, when such properties are merely the semantic shadows of the fact that they function as predicates.”⁸

I myself believe that normative statements can be true, can be facts in this minimal sense, and that this is all we need. But questions remain about how this can be so. In addition to questions, which I will address in later lectures, about how we could discover such truths and about the practical significance of normative beliefs, there are questions about how normative truths are related to facts about the natural world. To address these, it will be helpful to begin by saying something about the “fact/value distinction.”

It is widely believed, by both realists and non-realists about the normative, that there is an important distinction, sometimes called an “unbridgeable gap,” between

⁷ *Thinking How to Live*, pp. 182-183.

⁸ “How to be an Ethical Anti-Realist,” in *Essays in Quasi-Realism*, p. 181.

“facts” and “values.” But the idea that there is such a distinction has been challenged in a number of ways.

This idea is often expressed as the thesis that no “value” statement is entailed by any set of “factual” statements. But, as A. N. Prior pointed out, this leads to difficulties about how these two classes of statements are to be understood.⁹ Let F be a statement agreed to be factual, and V any statement agreed to be one about value. What, then, about $F \vee V$? It follows logically from F. So if no value judgment can follow from a factual one then it must be factual. But from $F \vee V$ and $\neg F$ one can deduce V. So $F \vee V$ cannot be factual if the thesis of non-derivability holds.

The existence of “thick” ethical concepts such as “cruel” is also often cited as evidence for the thesis that facts and values are inextricably intertwined.¹⁰ On the one hand, certain psychological characteristics seem sufficient to make a person cruel. But, on the other hand, calling someone cruel is clearly a value judgment, a form of moral criticism. Moreover, it does not seem possible to factor the meaning of ‘cruel’ into factual and moral components, since one cannot understand the factual component (which psychological traits it is that make someone cruel) without understanding the ethical point that makes the charge of cruelty a form of criticism. This point is generally made with reference to ethical concepts, but it is plausible to suppose that there are “thick” non-moral normative concepts that have analogous properties. “Unreasonable” and “closed minded” come to mind as possible examples.

In addition to these two problems for the idea of a fact/value distinction, it is sometimes pointed out that the justification for factual conclusions—scientific

⁹ A. N. Prior, *Logic and the Basis of Ethics*??

¹⁰ Hilary Putnam, ----; Bernard Williams, *Ethics and the Limits of Philosophy*

conclusions, for example—often involves appeal to what are clearly values, such as simplicity, clarity and the like. So some paramount factual statements depend on claims about value.¹¹

Despite these points, which certainly have some validity, there seems to be an important distinction between normative and non-normative claims. The question is what this distinction is. The first step to answering this question is to identify the classes of statements that are being distinguished.

As Hilary Putnam has observed, claims about the gap between facts and values have generally taken as their starting point some definite characterization of factual statements, with which value judgments are then contrasted. For Hume, ‘is’ statements were identified with “relations of ideas” and “matters of fact.” For the logical positivists, the relevant class consisted of analytic statements and those that are empirically verifiable.¹² I propose to start on the other side, beginning with a characterization of a class of normative statements. My basic idea is that normative statements are statements about the reasons that people have. But his thesis will require some refinement.

I have said that the characteristic element in normative judgments is the relation $R(p, c, a)$: “p is a reason for a person in situation c to do a.” But p cannot be a reason unless it is the case that p. So it would seem to follow the non-normative claim that p is not the case that the normative claim $R(p, c, a)$ is false. This might seem to be a case of inferring a normative conclusion from purely non-normative premises. The essentially normative content of R, however, is independent of whether p is true: it lies in the claim that, whether p is the case or not, if p *were* the case it would be a reason for someone in c

¹¹ See Putnam, *Op. cit.*

¹² Putnam, ---

to do a. So I will understand R in this subjunctive form and take what I will call a *pure normative claim* to be a claim of the form $R(p, c, a)$, where R is understood in this way.

As a first step, I will take the thesis of the autonomy of the normative from the naturalistic to be the thesis that no pure normative claim is entailed by any combination claims about physical and psychological facts. Given any combination of naturalistic claims, it is a further claim that some pure normative claim is true.¹³

Most of what we commonly think of as normative claims are not pure normative claims, but *mixed normative claims*. They involve pure normative claims but also make or presuppose claims about natural facts. “p is a reason for x to do a” is a mixed claim, since cannot be true unless p is true, and it also presupposes that there are conditions c such that the pure normative claim $R(p, c, a)$ is true and x is in circumstances c.

The idea that there is a “logical gap” between normative and non-normative claims may itself seem puzzling, since we commonly make what may seem to be sound inferences from non-normative facts to normative conclusions. For example, from

(1) If Jones does not leave the burning building now, he will be killed.

it seems to follow that

(2) Jones has reason to leave the burning building now.

If there is a logical gap between the normative and the non-normative, how can it be that we leap over this gap with ease many times every day?

The answer is that the “gap” consists in a failure of logical entailment, and the sense in which (2) obviously “follows from” (1) is not logical entailment. Rather, (2) seems obviously to “follow from” (1) because we take it to be obvious that

¹³ As Gibbard puts it, it is coherent to affirm any such combination of physical and psychological claims while denying $R(p, c, a)$. *Thinking How to Live*, p. 25.

(3) Jones's situation is such that the fact that doing a is necessary for him to avoid dying now is a reason for him to do a.

(3) is still a mixed normative claim, since it involves a claim about what Jones's situation actually is. But we could put these conditions c, whatever they are (facts about Jones' life in virtue of which he has reason to want to go on living) into the earlier premises. There would still be a "gap" represented by the pure normative claim that if someone is in these circumstances then he has reason to do what is necessary to prolong his life.

Since, in order to "get from" non-normative claims to some normative claim we need to make a claim about reasons, such a move will in general, as in this particular case, depend on some pure normative claim or claims. It is therefore unsurprising that one cannot "get from" any conjunction of non-normative premises to a normative conclusion without "already making" some normative claim. This reflects no infirmity or problem about the status of normative claims. It is a reflection simply of the kind of thing that pure normative claims are: they assign normative significance to certain non-normative claims.

The distinction between normative and non-normative claims is most likely to seem like a "gap" that it is difficult to get across if we focus on mixed normative claims, such as (2), in which the relational character of pure normative claims is not apparent.. The same is true of other normative claims that are often mentioned in this context, such as claims that something is good or is morally wrong. These claims appear simply to assign to their subject some normative property, and the gap is between having this property and having various natural properties. The relation between normative and non-normative is clearer when we focus instead on pure normative claims which have exactly

the function of assigning normative significance to the non-normative. This assignment of significance is just what any move across the supposed “gap” involves. So it should not be surprising that one cannot make this move without “already” making a normative claim.

Stating the thesis of the autonomy of the normative in this way allows us to capture what seems intuitively correct about the fact/value distinction while avoiding the problems raised by Prior. It applies in the first instance only to pure normative claims, and leaves open the question of mixed statements of the kind that Prior mentioned. It also enables us to accommodate the objections raised by Putnam.

The claim that we have reason to believe a particular empirical proposition is a normative claim. This is so even if the reasons cited are purely epistemic: considerations that count in favor of the truth of the proposition in question. Beyond this, our reasons for accepting a scientific theory may, or may not, be based on further reasons that are not all truth-related.¹⁴ But neither of these possibilities indicates a puzzling intermingling of facts and values. If they appear to do so this is due to a failure to distinguish between our reasons for accepting a theory, of for counting a proposition true, and the claims made by that theory or proposition. The former may be normative; the latter is not.

Nor is there a troubling intermingling of fact and value in “thick ethical concepts.” To claim that Caligula was cruel is certainly to make a claim about what he *saw* as a reason and responded to in his actions, and what, on the other hand, he was generally indifferent to. Such claims attribute normative views to Caligula, but do not

¹⁴ For discussion, see Joseph Raz, “Reasons: Practical and Adaptive.”

make any such claims themselves. The claims made, so far at least, are purely psychological. But genuinely normative elements may enter in two related ways.

A normative element enters insofar as the charge of cruelty involves not only the claim that Caligula was indifferent to certain concerns, such as the suffering of his victims, but also the claim that these were things that he should not have been indifferent to, because they really are reasons. A second normative element lies in the idea that cruelty is something one has reason to condemn and more generally, that one has good reason to react differently to someone who is cruel than to someone who is not—to avoid their company, not to trust them, and so on. These two elements are related. The reasons one has to respond differently to someone who is cruel depend on the particular reasons that a cruel person is insensitive to, and the importance of responding correctly to these reasons.

These two normative elements are central to the concept of cruelty. They give the content its point and unify its empirical content. It is questionable whether a person who failed to understand these normative elements could grasp the concept and readily tell which psychological traits and forms of behavior count as cruel. But it would be odd for even someone who understood these normative elements to *use* the concept unless they shared the normative judgments that they involve. Oscar Wilde, for example, understood the normative elements in the concept of blasphemy, but he said that ‘blasphemy’ was “not a word of his,” because he did not share these normative views.

It may be a complex task of analysis to identify the particular empirical and normative elements in a particular thick normative concept, and to determine which normative elements are asserted when that concept is claimed to apply and which are

only presupposed. The purpose of the sketch I have given is merely to explain why such concepts do not undermine the thesis that there is an important distinction between normative and non-normative claims. That thesis is not that every statement falls into one or the other of these two classes. It claims only that there is a class of normative concepts, and statements using them, that cannot be analyzed in terms of statements using only certain other concepts.

What I have tried to do so far is to explain why pure normative claims are not entailed by non-normative claims. But it is widely agreed that the normative nonetheless supervenes on the non-normative. This relation has two aspects. First, normative facts can depend on certain non-normative facts: they vary when these non-normative facts vary. Second, normative facts are fixed by the non-normative facts: they cannot vary when non-normative facts *do not* vary. It has been held to be a problem for views that allow for the existence of irreducibly normative truths to explain why they are related to non-normative truths in this way.¹⁵

This phenomenon might be seen as a kind of metaphysical necessity—the (metaphysical) impossibility of a world that is like ours in all non-normative respects but in which different normative facts obtain. But this is not an accurate description of the phenomenon in question. The necessity involved in the supervenience of the normative on the non-normative is not metaphysical but normative.

To understand the phenomenon of supervenience it is important to be clear what kind of normative claims are involved. The normative facts that can vary as non-normative facts vary are facts that consist in the truth of mixed normative claims, such as

¹⁵ See Mackie, *Ethics: Inventing Right and Wrong*, p. 41; Blackburn, *Essays in Quasi-Realism*, pp. 114-148; Frank Jackson, *From Metaphysics to Ethics*, pp. 118-125.

the claim that someone has a reason to do a certain action, or that a particular consideration is such a reason. In order for some fact to be a reason, it has to be a fact, and even if it is a fact its being a reason may depend on other facts.

So, for example, the fact that it would be very painful to put my hand into a flame is a reason not to do so. But if putting one's hand into a flame was not painful, then "the fact that it would be very painful to put my hand in to a flame" would not be a fact, and I would not have the reason just mentioned. So normative facts depend on non-normative ones, and which ones they depend on is a normative matter, specified by normative facts consisting in the truth of pure normative claims.¹⁶ The truth of pure normative claims, by contrast, does not depend on, or covary with, non-normative facts.

The further question is whether the truth of pure normative claims can vary "on its own," not as a result of variation of other facts on which these claims depend. The answer is that the truth of pure normative claims does not vary in this way, and the fact that it does not (that these claims are not in this way "contingent") is a normative matter, not a metaphysical one. Perhaps putting one's hand into a flame might not have been painful, or harmful. If so, then the fact that I have reason not to put my hand into a flame is also contingent: I would not have had such a reason if putting my hand into the flame were neither harmful nor painful. But the fact that if putting my hand into the flame would be painful, this would be a reason for me not to do so is not something that "might have

¹⁶ Blackburn appears to agree. He writes, "A quasi-realist will see both covariance and the asymmetry of dependency as a reflection of the fact that valuing is *to be done* in the light of an object's natural properties, and without that constraint nothing recognizably ethical could be approached at all." "Supervenience Revisited," *Essays in Quasi-Realism*, p. 146.

been otherwise.”¹⁷ I confess that I do not know how to argue for this claim. It seems to me evident from reflection on what basic normative truths are. But it does not seem to me, on reflection, to be something that we should find puzzling. Pure normative truths are not contingent in the most obvious way—that is, dependent on contingent facts about the natural world. Given that they are not contingent in this way, why should we expect them to be contingent in some further sense?¹⁸

The autonomy of the normative as I have formulated it is a thesis about the content of normative concepts and, hence, the content of normative claims. But many who accept the idea that normative concepts are non-natural maintain that things are different when it comes to normative *properties*. Allan Gibbard and Mark Schroeder, for example, agree that normative concepts cannot be analyzed in naturalistic terms, but they maintain, in different ways, that the properties signified by normative terms are

¹⁷ Kit Fine argues for the thesis that the necessity of normative claims is distinct from and not reducible to metaphysical necessity in “The Varieties of Necessity,” Gendler, Gendler-Szabo and Hawthorne, eds., *Conceivability and Possibility* (Oxford: Oxford University Press, 2002), pp. 253-281.

¹⁸ If this is correct, then it seems to follow that the supervenience of the normative on the non-normative is quite different from other kinds of supervenience, such as the supervenience of the mental on the physical. Normative judgments make claims about (the normative significance of) non-normative facts. That normative facts covary with non-normative facts, and the particular facts on which they depend, is determined by normative truths, and that normative truths do not vary “on their own” is also a normative matter. By contrast, it is not part of the *content* of claims about mental phenomena that they attribute mental properties to physical states, and “mentalistic” truths do not specify which physical states mental states depend on or covary with. Confidence that the mental covaries with the physical and cannot vary independently arises rather from the acceptance of what might be called the hegemony of the physical: the thesis that all natural phenomena have physical explanations.

naturalistic.¹⁹ So I need to consider the relation between concepts and properties, and how normative properties should be understood.

In one sense, the distinction between concepts and properties is clear. Identifying concepts is a matter of determining the content of our thoughts. Specifying properties is a matter of determining the nature of things in the world to which those concepts apply. The question is when and why a characterization of the property corresponding to a concept will go beyond what is specified by that concept itself. According to what might be called a minimal understanding of a property, the property signified by a concept is just a matter of having those features included in the concept. So if one understands the concept, then there is no more to be said about the property. In some cases, however, this view seems mistaken: there is more to be said about what it is to be a thing in the world of the kind to which the concept applies. The interesting question is when and why this is the case.

Consider first a naturalistic concept, such as *water*. If the concept, water, is defined only by features of water that figure in our everyday experience, such as “colorless liquid that falls as rain and fills rivers, streams, lakes and oceans,” then there is more to be said about what water is: for example, that it has the chemical composition, H₂O. We might say, then, that the property of being water is a matter of having those physical characteristics, whatever they may be, that are responsible for its having the observed characteristics that figure in the concept. Something similar might be said about other concepts of natural kinds, such as the concept of lightning. In all these cases, the fact that there may be more to the property signified by a concept than is specified in that

¹⁹ In different ways because Schroeder is a reductive naturalist while Gibbard is an expressivist. See *Thinking How to Live*, pp. xxx and *Slaves of the Passions*, pp. sss.

concept is not a metaphysical fact but a reflection of the fact that the concepts in question identify *natural* phenomena on the basis of certain features that are apparent to us. This leaves open the possibility that there is more to be said about the nature of these phenomena—that is to say about what in “the world” described by chemistry and physics is responsible for these features.

There may be a broad parallel in the case of some normative concepts. The concept, morally wrong, for example, includes such things as “action that anyone has very strong reason not to perform, and which makes guilt appropriate on the part of one who has done it and resentment on the part of those to whom it is done.” But this leaves open what reason there is not to perform these actions, and to feel guilt and resentment as a result of their being done. A person can understand and employ the concept *morally wrong* without having a clear idea what these reasons are, just as someone can have the concept, water, without knowing its chemical composition. There is more to be said about what makes something morally wrong, and the task of giving this further account might be said to be the task of characterizing the property of moral wrongness.²⁰ If this is right, then this is another case in which the minimal understanding of a property is insufficient.

It is worth emphasizing that neither in this case nor in that of natural kinds such as water is the further account that can be given of a property, beyond what is contained in the corresponding concept, a metaphysical account. The further characterization of a natural kind such as water is scientific—a matter of chemistry and physics. The further

²⁰ I took this line in *What We Owe to Each Other* (see p. 12). Derek Parfit criticized it in “Justifiability to Each Other,” *On What We Owe to Each Other*, Philip Stratton-Lake, ed. (Oxford: Blackwell Publishing, 2004) pp. 67-89. For further discussion, see my “Wrongness and Reasons: A Reexamination,” in Russ Shafer-Landau, ed., *Oxford Studies in Metaethics, Vol 2* (Oxford: Oxford University Press, 2007), pp. 5-20.

explanation of moral wrongness is normative—a matter of identifying the relevant reasons.

My concern in this lecture is with the concept of a reason—more exactly, the relational concept $R(p, c, a)$. So the question before us is whether there is something further to be said about what it is to be a reason, beyond what is given just by this relational concept, something further that might be said to identify the property signified by that concept. Since the concept of a reason, like that of moral wrongness, is a normative concept, it would seem that any further characterization of what it is to fall under that concept would also need to be normative. But if the domain of the normative consists solely of claims about reasons, then no normative characterization of the concept of a reason itself can be given, since it would have to employ this very concept. So, given this way of understanding normativity, and this idea of what a further explanation of the concept would have to be like, it seems to follow that the concept of a reason is fundamental.

The possibility of a concept that is fundamental in this way cannot be ruled out. Suppose that the property of being water is that of being H_2O . What then should we say about property signified by the concept H_2O ? Perhaps this concept is fundamental: the property of being H_2O is just that, being H_2O . But perhaps there is some further characterization of what it is to have that molecular structure. If so, then we can ask again, about the property signified by that concept. Presumably at some point we reach a concept that is fundamental: the property it signifies cannot be characterized in any deeper way than by that concept itself.

The argument I sketched for thinking that the concept of a reason is fundamental depended on two claims, both of which may be questioned. The first was that normativity is to be understood solely in terms of reasons. I will return to this question in a later lecture. The second was that if some further account of the (relational) concept of being a reason were to be given, this account itself would (as in the case of moral wrongness) have to be normative and therefore (it was assumed) have to employ the concept of a reason in a way that would render it uninformative. This might be questioned in two ways. Perhaps what is needed is not a normative account of reasons but an explanation of some other kind. Or perhaps a normative account could employ the concept of a reason but nonetheless be informative, by, for example, characterizing the domain of reasons in a helpful way. I will explore the first of these possibilities here, returning to the second later on.

One suggestion would be simply that it is metaphysically puzzling what the property corresponding to the relational concept of a reason could be, and that this calls for some kind of (non-normative) explanation. Allan Gibbard, for example, although accepts the idea that there are normative concepts, objects to the idea that there are normative properties, which he would find metaphysically odd.²¹ He does not say exactly why, but it may derive part of its plausibility from the ambiguity I have mentioned about the idea of “the world.” If by a normative property one meant a property in the physical world, then I agree that this would be odd. But if all one means by something’s having the property of being a reason is that the concept, reason, properly applies to it, then I see

²¹ *Thinking How to Live*, p. 32.

no oddity. There would be a normative property only in a minimal sense corresponding to the minimal notion of truth, which Gibbard seems to accept.

I believe that it is a mistake to think that there is a *metaphysical* puzzle here. As I noted above, the kind of explanatory need that drives us to a further account of the property of a natural kind, or of moral wrongness, is not metaphysical, but scientific in the first case and normative in the second. Nonetheless, we should consider some candidate accounts of the property of being a reason, and see what kind of explanation they might provide.

If there is some more substantive account of the property of being a reason, what might it be? Gibbard, discussing the normative concept *good*, says that the property signified by this concept is not the *de facto* extension of that concept but rather those things that are necessarily good. I take this to mean those things that are or would be good, contingencies aside. Transferring this to the three place relation of “being a reason for,” the suggestion would be that the property signified by this relation would be all the triples $\langle p, c, a \rangle$ such that p would be a reason for a person in situation c to hold attitude a if p were the case and there actually was a person in this situation. That there is such a set, or collection, seems unproblematic. But it does not seem a plausible candidate for the role in question. This collection would tell us which things are, or would be reasons for other things in certain circumstances. But it would not explain what it is to be a reason or what makes something a reason. It thus stands in a very different relation to the concept of a reason than, say, having the chemical composition H_2O stands in to the concept water.

In this respect there is more to be said for the naturalistic account of the property of being a reason offered by Mark Schroeder. Like Gibbard, Schroeder agrees that the concept of a reason cannot be analyzed in non-normative terms.²² But he believes that the property of being a reason can be so analyzed. Specifically, he believes that for p to be a reason for a person in situation c to do a is for there to be some q such that that person has a desire whose object is q and the truth of p is part of what explains how that person's doing a promotes q.²³

This account of what it is to be a reason is stated in purely naturalistic terms. So it may seem doomed at the start. To identify being a reason with a naturalistic property seems immediately to destroy its normativity. Schroeder's response that if normativity consists in being properly analyzed in terms of reasons (as I would agree), then his account preserves normativity as long as it captures the idea of a reason. One may think (as I do) that its very naturalistic character prevents it from doing this. But beyond merely asserting that his naturalistic account cannot preserve the normative character of reasons, we should consider Schroeder's reasons for believing that it does.

What would be required in order for this analysis to be successful? First, there would have to be a reasonably good extensional fit with our firmest intuitive judgments about reasons.²⁴ This is not sufficient, however. The account proposed above following Gibbard's model, for example, would meet this condition. But it seemed a poor candidate

²² *Slaves of the Passions*, p. 65.

²³ *Ibid.*, p. 59. I have modified Schroeder's definition slightly to fit my statement of R(p, c, a). One apparent difference, which I will set aside for the moment, is that his definition, on the face of it, applies only to reasons for action.

²⁴ Schroeder argues, in Chapters 5 and 6 of *Slaves of the Passions*, that his account meets this condition. I do not find those arguments persuasive, but I will leave this disagreement aside.

to be the property of being a reason in part because it did not thing to explain the features that reasons have. Schroeder believes that his account does just this. Specifically, he believes that it explains why someone is motivated by the belief that he or she has a reason to do something, explains why facts about reasons supervene on natural facts, and explains why the reasons that some people have differ from the reasons that others have. I agree that Schroeder's account offers explanations of the first two kinds, although I believe that a non-reductive account can provide explanations that are equally good, if not better. I have already explained how a non-reductive account can explain supervenience, and I will return to the question of motivation in a later lecture. So I will focus here on the third claim.

Schroeder's main example, which he returns to throughout the book, involves two people, Ronnie and Bradley. Both have been invited to a party where there will be dancing. "But," Schroeder says, "while Ronnie loves to dance, Bradley can't stand it." He claims, plausibly, that the fact that there will be dancing at the party is a reason for Ronnie to go to the party but not a reason for Bradley to go. Moreover, it seems uncontroversial that this difference between Ronnie's reasons and Bradley's is explained by "some feature of their psychologies." The Humean Theory of Reasons, as Schroeder understands it, is that "Every reason is explained by the kind of psychological state that explains Ronnie's reason in the same way as Ronnie's is."²⁵ (Schroeder sees his own view as one particular version of The Humean Theory.)

If what is to be explained is the difference between Ronnie's reasons and Bradley's then it does seem uncontroversial that this difference lies in something about

²⁵ *Slaves of the Passions*, p. 2

their psychologies. But this is more controversial if what is in question is the (most fundamental) explanation of Ronnie's reason to go to the dance. It is very plausible to say that what explains the difference between Ronnie's reasons and Bradley's is the fact that Ronnie enjoys dancing and Bradley does not.²⁶ But this leaves open the question of why the fact that Ronnie enjoys dancing makes it the case that the fact that there will be dancing at the party gives him a reason to go? This might be, as Humean theories hold, because Ronnie has a desire for experiences that he finds pleasant. Or it might be, as many non-Humean theories would maintain, simply because people have reason to do what they find pleasant.²⁷ So, although it may be non-controversial that what explains the difference between Ronnie's reasons and Bradley's is something about their psychological states, it is controversial whether the most fundamental explanation of Ronnie's reason is a psychological state. Indeed, this is just the point at issue between Humeans and non-Humeans.

The possibility of a hedonistic explanation of Ronnie's reason for going to the dance comes up at two further points in Schroeder's argument. The first is in his interesting discussion of what he calls the "no background conditions view." This is the view that any condition that is needed in a full explanation of why something is a reason

²⁶ Schroeder considers this possibility as one candidate for the psychological feature that, according to a Humean Theory, explains Ronnie's reason, hence as one possible variant of a Humean view. (3) What I am suggesting is that this explanation of Ronnie's reason could be offered by a non-Humean theory, and that this possibility undermines the support that the example of Ronnie and Bradley offers for a Humean theory.

²⁷ There is also a question here about time. Is the psychological state that explains Ronnie's reason a state that he is in at the time he is deciding whether to go to the party, such as the fact that he wants to dance at that later time, or a desire for experiences that, at the time of their occurrence, he will find enjoyable? Or is the fact that he has reason to go to the dance explained by a future psychological state, the pleasure that (he has good reason to expect) he will feel when dancing at the party?

for a person reason to perform given action must itself be part of that reason. (23) If this view were correct, then on a Humean theory a full statement of every reason for action would make reference to the agent's desires. This would, Schroeder says, give all reasons an implausible self-regarding character, suggesting that all agents are ultimately moved only by the satisfaction of their own desires. Since not all reasons seem to have this self-regarding character, this would count against the plausibility of Humean theories.

Schroeder's response is to argue that not every factor that is needed to explain why a certain consideration is a reason for an agent is also part of that reason. "If Ronnie genuinely desires to dance, then *all it should take* for him to be moved to go to the party is the thought that there will be dancing there."(27) There is no need for him to think also "and I desire to dance." This general point, about the distinction between reasons and background conditions, is quite correct, and important. It is recognized in my formulation of the relation "being a reason for" by the distinction between p, which is the agent's reason for a, and those features of the agent's situation c in virtue of which p is a reason. But the application of this distinction to the case of Ronnie and Bradley seems to count against Schroeder's view rather than to support it.

In general, including "and I desire X" in the content of a reason gives the agent's action an *implausible* self-regarding character because in many cases the agent desires the thing in question for some reason not connected with the satisfaction of his or her desires. If a person desires to contribute to the alleviation of world poverty, it is implausible to say that part of her reason for sending a check is that this will fulfill her desire. But if Ronnie goes to the party because he likes to dance, then his reason for going is most plausibly understood as having a self-regarding character that it would not have if, for

example, he desired to go to the dance because he had promised to take his girlfriend dancing (even though he did not much enjoy it himself), or if he desired to go, and to dance, because he wanted to encourage his younger siblings' interest in dancing, in order to keep them from more dangerous pursuits. In cases of the latter kinds, including the fact that he desires to accomplish the further end in question as part of Ronnie's reason would give that reason an implausibly self-regarding character even if it were true, as Schroeder maintains, that such a desire was a necessary condition of those ends being reason-providing. But Ronnie's reason in the case as Schroeder describes it *is* self-regarding. This suggests to me that the fact that Ronnie "likes to dance" plays a different role in that case than the general role that desire would play in the other cases I have mentioned if Schroeder's view were correct. What it suggests is that the psychological state that differentiates Ronnie's situation from Bradley's is not a desire (playing the same role as desires in these other cases) but rather the fact that Ronnie enjoys dancing, and that this fact is part of Ronnie's (unobjectionably self-regarding) reason, not just a background condition, as desires may be in these other cases. It is, of course, a further question, and a matter in dispute, whether the fact that he enjoys dancing provides Ronnie with a reason to go to the party only given the background condition that he desires pleasant experiences.

These issues arise again at the beginning of Schroeder's Chapter 8, where he briefly considers desires and "what people take pleasure in" as alternative candidates for the role of "the psychological state ... which most fundamentally explains the difference between Ronnie's and Bradley's reasons." (146) One way to decide between these alternatives would be to imagine cases in which Ronnie would take pleasure in dancing

but does not know this and has no desire to dance, and to consider what reason he would have to go to the party if this were the case. Schroeder rejects this method of argument, on the ground that our intuitions about such cases are unreliable. He says that we can't, for example, rule out the possibility that even if Ronnie has no desire to dance his reason for going to the party depends on some other desire, such as a desire to enjoy himself. (147) It would be difficult, Schroeder says, to screen out the possibility of such a desire, or the possibility that Ronnie has some other desire that explains a reason for Ronnie to do what he enjoys. "So it seems more promising," he says, to proceed instead by "taking a closer look at what kind of psychological state is most *suited* to explain the existence of reasons" subject to constraints he has outlined earlier.

We should note two things about this move. First, it seems extremely plausible that Bradley, as well as Ronnie, desires to do what he enjoys, or that he has some other desire that explains why he has reason to do such things. It is therefore very plausible to suppose that the difference between Bradley and Ronnie lies somewhere else, such as in facts about *what* they enjoy. Second, Schroeder's strategy seems to involve a shift away from looking for an explanation of the difference between Ronnie's and Bradley's reasons to looking instead for a kind of psychological state that is suited to explain the existence of reasons in general, and Ronnie's reasons in particular. But, as I have said before, the idea that it is a psychological state that we should be looking for was made plausible to begin with by the fact that we were looking for an explanation of the difference between Ronnie's and Bradley's reasons. Although seems very plausible that this difference must lie in their psychological states, the idea that all reasons are

explained by psychological states is a different, and much more controversial claim, not obviously supported by the example of Ronnie and Bradley.

Frank Jackson has also objected strongly to the idea that there might be normative properties in addition to the purely naturalistic properties with which they are co-extensive. It is possible that he would not consider these objections to apply to normative properties understood in the minimal way I have proposed. But I should consider whether the objections he raises apply to my proposal. In the terms we have discussed, what Jackson is opposed to is taking the property signified by R to be something other than the set of triples $\langle p, c, a \rangle$ that Gibbard identified as the naturalistic property signified by this relation.

Jackson mentions three objections. The first is that “it is hard to see how we could ever be justified in interpreting a language user’s use of, say, ‘right’ as picking out a property distinct from that which the relevant purely descriptive predicates pick out, for we know that the complete story about how and when the language user produces the word ‘right’ can be given descriptively.”²⁸ Suppose we know the set of triples $\langle p, c, a \rangle$ such that a language user assents to $R(p, c, a)$. Does this amount to “the complete story” about how that language user understands the relation R? It seems to me that it does not. What we need to know further is how that language user responds when he believes that the relation $R(p, c, a)$ holds. In order to know whether the language user assents to $R(p, c, a)$ just when he or she takes it to be a “true thought” that p counts in favor of a for someone in c, we need to know whether he or she generally treats $R(p, c, a)$ as relevant to

²⁸ Jackson, *From Metaphysics to Ethics* (Oxford: Oxford University Press, 1998), p. 127.

the question of whether to do a when he or she takes him or herself to be in circumstances c and believes p.

This also provides a basis for responding to Jackson's second objection, which is that "it is hard to see how the further properties could be of any ethical significance. Are we supposed to take seriously someone who says, 'I see that this action will kill many and save no one, but that is not enough to justify my not doing it; what really matters is that the action has an extra property such that only ethical terms are suited to pick out'?" In short, the extra properties would be ethical idlers."²⁹ The property minimally signified by R, on my view, is not a "normative idler." To claim that R(p, c, a) holds is precisely to claim that p justifies a; it is not to claim that p has some further property which does the justifying.

Finally, Jackson asks how we determine in which cases there is, in addition to some purely descriptive property, a normative property coextensive with it. The answer is that this is in each case a normative question: it depends on whether a particular p actually is a reason for someone in c to do a.

I have been defending the idea that there are irreducibly normative truths about reasons. This claim of truth is extremely minimal. In particular, I am not claiming that there is a (relational) property "in the natural world" corresponding the (relational) concept "being a reason for." Normative truths, in my view, constitute a distinct realm and need no metaphysical reality in order to have the significance that we commonly grant them.

²⁹ *Ibid.*

Given the minimal nature of my claims of truth for normative assertions, it may be asked how much my view really differs from Gibbard's expressivism or Blackburn's quasi-realism. Both Gibbard and Blackburn allow for, or even embrace, the idea of normative claims being true in a minimal sense. And like them, I am claiming that normative judgments are about our reactions to the natural world, rather than about that world itself (specifically, in my case, about the appropriateness of these reactions.) So it may seem that little difference remains. As a challenge to my view, this would be the correlate to challenges that have been made to Blackburn, that his quasi-realism was no different from realism.³⁰

Despite these appearances, important differences remain. They have to do with the way in which the practical significance of normative commitments is explained, with the way in which interpersonal advice and disagreement about normative questions is interpreted, and with the sense in which the correctness of our normative commitments is independent of those commitments themselves. I will discuss these matters in the next lecture.

³⁰ See Gideon Rosen, ----, Ronald Dworkin, ----

