

The John Locke Lectures 2009

Being Realistic about Reasons

T. M. Scanlon

Lecture 2: Metaphysical Objections¹

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.

¹ For comments that have helped me in revising this lecture I am indebted to Robert Adams, Anthony Brueckner, Tyler Burge, Matti Eklund, Tom Hurka, Derek Parfit, Charles Parsons, Wlodek Rabinowicz, Michael Ridge, Ralph Wedgwood, and Tim Williamson.

Stating his version of this objection, John Mackie writes that, “if there were objective values, then they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe.”² Others have made similar objections. It is natural to describe this as an ontological objection: that the idea that there are irreducibly normative truths has implications that are incompatible with plausible views about “what there is.”

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 a “universe of discourse” 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

² *Ethics: Inventing Right and Wrong*, p. 38. 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, 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 of any kind, not just objective moral values. At least this is how I am going to take his argument, I hope not unfairly.

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 ontology we have reason to accept. 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 thinking about reasons—and simply accept 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.

An ontological objection to normative truths depends on some restrictive view of this kind. Mackie’s objection seems to be based on the view that all of our ontological commitments must be claims about what exists in the physical world of space and time. This is what is suggested by his remark that objective values would involve entities, qualities or relations “different from anything else *in the universe*.” And the same idea is reflected in the charge that the idea that there are irreducibly normative truths would be incompatible with a scientific view of the world. In each case, ‘the universe’ and ‘the world’ seem to refer to the physical universe of particles, planets and so on, that is described by science. The idea that our ontological commitments should be restricted to

things in this world may strike many as a sensible naturalism. But I believe that it is a view we should not accept.

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. 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 a scientific view of the world does not mean accepting the view that this world and the things in it are the only things we can refer to and talk about sensibly, or have true beliefs about, or the only things that we should be ontologically committed to in Quine's sense (that is, to “quantify over.”)

In Quine's case, the identification of things we have reason to quantify over with things we take to exist in the spatio-temporal world is built in from the start, because the theories whose ontological commitments he is concerned with are theories of that world: accounts of the world that impinges on our sensory surfaces. This immediately excludes the normative, absent some naturalistic reduction. It might seem also to exclude ontological commitments to mathematical entities such as numbers and sets (absent such a reduction.) But this is not necessarily the case. Quine's holism treats mathematical and logical truths as the most abstract parts of our theory of the world, which faces the tribunal of sensory experience as a whole. Whether the best such theory should quantify over numbers or sets is a scientific question about what the best (most successful and simplest) overall scientific theory is like. Thus, speaking of the suggestion by some logicians that “all the mathematical needs of science can be supplied on the meager basis of what has come to be known as predicative set theory,” Quine writes that “Such gains

are of a piece with the simplifications and economies that are hailed as progress within natural science itself. It is a matter of tightening and streamlining our global system of the world.”³

It does sound odd to say that *the world contains*, in addition to particles and mountains and planets, numbers, sets, and perhaps rights and obligations as well. This sounds odd, even to someone like me who believes that there are such things. It sounds odd because, taken as a claim about physical reality, it seems obviously false. I believe that the way of thinking about these matters that makes most sense is what I called above a permissive first-order view, which does not privilege science, but takes as basic a range of first-order domains or practices, such as mathematics, science, and moral and practical reasoning.

A *domain*, in the sense I have in mind, 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 these concepts; and some accepted procedures for settling questions employing these concepts. What these procedures are is often vague or incomplete, and there is room for argument, internal to the subject in question, about how they are best understood: mathematical disagreement about standards of proof in mathematics, scientific disagreement about the adequacy of certain kinds of evidence, moral disagreement about criteria of right and wrong, and so on. In contrast to Carnap, whose view is similar to mine in some respects, I do not take the procedures appropriate to a domain to be determined by “linguistic rules” for the use of

³ *Pursuit of Truth* (Cambridge, MA: Harvard University Press, 1990) p. 95.

the terms in question.⁴ People can use terms such as ‘number,’ ‘set,’ or ‘wrongness’ without misuse of language while disagreeing to some degree about the facts about such things, and about the best ways of determining these facts

On the view I am proposing, questions about the existence of entities of a given kind are, in general, properly settled by the modes of reasoning of the domain to which they belong: questions about the existence of numbers settled by mathematical argument, questions about the existence of subatomic particles by applying scientific methods, and so on. This is not to say, that the standards actually accepted by practitioners in such a domain at a particular time have the last word about what is true about objects in that domain. Standards can evolve, as a result of further reflection about the subject matter in question.

Nor does my view entail that these first-order domains are entirely autonomous, and that nothing beyond the (evolving) criteria of a domain is relevant to the truth of statements within it. For one thing, statements in one domain can conflict with those in another, and when this happens these claims need to be reconciled, and some of them perhaps given up. We might, for example, 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

⁴ “Empiricism, Semantics and Ontology,” in *Meaning and Necessity: A Study in Semantics and Modal Logic* (Chicago: University of Chicago Press, 1956), pp. 205-221.

these conflicts 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.⁵

To put the same point more generally: there can be meaningful “external” questions about the adequacy of the generally accepted modes of reasoning in a domain, and about the truth of the statements, including existential statements, that these modes of reasoning support. Meaningful questions of this kind are questions about whether the accepted modes of reasoning in a domain actually support conclusions of the kind required in order for statements employing the concepts of that domain to be true and to have the significance that is claimed for them. They are external questions insofar as they cannot be settled by the modes of reasoning of the domain in question. But they are made meaningful by the claims made by and for statements internal to the first-order domain itself.

In some cases, as in the example of witches and spirits, these external questions are questions about the natural world, to be settled by scientific criteria. Some claims about gods may involve claims of this kind, such as claims about the creation of the universe, or about what is happening when lightning occurs. An interpretation of moral claims, or claims about reasons for action that takes them to be straightforwardly true could would be “incompatible with a scientific view of the world” if these claims involved, or if their supposed significance presupposed, claims about the natural world that science gives us good reasons to reject. But irreducibly normative statements do not

⁵ My view is thus not open to the objection that Hartry Field raises, unfairly I think, against Crispin Wright’s similar view, the objection that giving priority to “ordinary criteria” would commit us to the existence of God and to the idea that Zeus throws thunderbolts. See Field, “Platonism for Cheap? Crispin Wright on Frege’s Concept Principle,” in his *Realism, Mathematics and Modality* (Oxford: Basil Blackwell, 1991), p. 155.

involve or presuppose such claims. Nor, as I will argue in Lecture 4, does the possibility of our coming to know normative truths involve claims about causal interaction. Nor does the possibility of agents being motivated by their beliefs about normative facts involve causal interaction with these facts. It is true, however, that normative claims would not have the significance that we normally attribute to them if there were no rational agents. So the existence of such agents is a presupposition of the practical domain that could in principle be undermined by external argument. I do not believe that it is in fact undermined in this way, since I believe that rational agents are just a kind of natural organism, and that such things do exist.

Morality has presuppositions of a different kind. The claims that we make about moral right and wrong generally presuppose 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, and ordinary ways of arguing for moral conclusions, do not make clear what these reasons are. There is therefore a question, external to morality, whether there are such reasons, and 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 scientific or metaphysical (unless metaphysics is understood to include normative claims), but rather normative, hence “practical” in one sense of the term—a question about what we have reason to do.

Our ontological commitments in Quine’s sense are determined by the totality of statements we accept: they consist of those things that must be included in the universe of discourse of a model that makes all these statements come out true. On the view I am proposing, we should decide what ontological commitments to have simply by applying

the relevant first-order criteria, taking into account the interaction between different first-order domains of the kind I have just been describing. So while there are, on my view, meaningful external questions about these claims, we have no reasons other than those arising from these various first-order domains to be concerned in general with what our ontological commitments are: no domain-independent reason to want to minimize these commitments, for example, and no reason in general to want limit these commitments to concrete entities as opposed to abstract ones. There may be good reasons in some cases for preferring simpler or more economical first-order 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 such preferences, and what counts as “simplicity” in the relevant sense, will be specific to those domains, not reflective of a general reason to prefer ontological minimalism.

The “universe of discourse” comprising our ontological commitments as a whole is a purely formal notion. Our having these commitments does not represent a claim on our part about what *the world* contains, in any meaningful sense of ‘the world.’ To say that it does invites, first, worries like Mackie’s, which arise from taking these to be commitments about the natural world, i.e. the physical universe.⁶ And if we respond to this first worry by denying that numbers, say, are part of the natural world, while still insisting that they are part of “the world” we invite questions about what this shadowy “world” is to which numbers and perhaps other non-spatial entities all belong. It is better

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.”

to avoid such questions altogether by making clear that the set of one's diverse ontological commitments is a purely formal notion and that there need be no world containing all these things.⁷

This move would be mere evasion if there were some general, domain-independent conditions of "existence" such that all of our various first-order existential claims entailed or presupposed that entities of the kinds they refer to fulfilled these conditions. If this were so, then there would be a genuine external question as to whether the things to which we are committed actually exist. But there are no such conditions.⁸ There is a purely formal idea of existence, captured by the logical rules governing the existential quantifier. These rules are the same across all domains. 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

⁷ Those holding positions more similar to the one I favor also 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." *Frege's Conception of Numbers as Objects* (Aberdeen University Press, 1983), p. 13. My view about the criteria of existence for numbers and other non-spatio-temporal objects is very close to Wright's. But saying that these things are "furnishings of the world" seems to me misleading, and to invite unnecessarily a response like Mackie's.

⁸ Here I agree with Hilary Putnam. See his *Ethics without Ontology* (Cambridge, MA: Harvard University Press, 2004), pp. 94-95, and with William Tait, who states a view about ontology very similar to the one I am advocating in "Truth and Proof: The Platonism of Mathematics," reprinted in his *The Provenance of Pure Reason: Essays in the Philosophy of Mathematics and Its History* (Oxford: Oxford University Press, 2005). See also Tait's remarks about ontology in his introduction to that volume, pp. 6-10.

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. Nothing more is claimed or required.⁹

One objection to the view I am recommending might be that it is too permissive. According to this view, it might be said, we could adopt some way of talking which specified criteria of identity for objects of a certain sort, and truth conditions for sentences containing terms referring to them, and which allowed for existential generalization from such sentences. According to my view, 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 existence of things quantified over in the existential statements counted as true in this way of talking. They would be among our “ontological commitments.” Can we take seriously an idea of existence that comes so cheaply?

A general version of this worry is what Richard Heck calls The Proliferation Problem.¹⁰ Heck is concerned with the version of the problem that is raised by the possibility of introducing new singular terms, and hence objects to which they refer, by

⁹ Does this mean that numbers, sets, obligations, and so on, “exist in a different sense,” that is to say, that the existential quantifier has a different meaning applies to numbers than it has when applied to elephants? We would not say that analogous thing about other logical operators, such as conjunction. It does not seem that the difference in the truth conditions of ‘Snow is white and grass is green’ and the truth conditions of ‘Two is even and the square root of two is irrational’ is even partly accounted for by differences in the meaning of ‘and’ in these two cases. Perhaps it is more plausible to say this in the case of ‘exists.’ But I want to set this question about meaning aside. What I am claiming is (1) that the only thing common to existential claims across all domains is the purely formal logic of the existential quantifier and (2) that the conditions required in order for objects in different domains to exist varies from domain to domain. Whether this variation is fully accounted for by the different sortal terms involved or also reflects variation in the meaning of ‘exists’ is a separate matter on which I take no position.

¹⁰ Richard Heck, “Syntactic Reductionism,” *Philosophica Mathematica* (3) Vol. 8 (2000), 124-149, p. 145???.)

means of abstraction on the model of Frege's remarks about "directions."¹¹ Frege proposed that we could introduce the term 'dir(a)' by means of the definition:

$$\text{dir}(a) = \text{dir}(b) \text{ iff } a \text{ is parallel to } b$$

(Frege also required, in order for 'dir(a)' to be a genuinely referring term, supporting existential generalization, that we must provide truth conditions for all sentences containing this term, not just for sentences of the form 'dir(a) = dir(b)'. I will return to this further requirement below.)

It may seem plausible that there are such things as directions. But Heck observes that we could introduce a new term in this way on the basis of any equivalence relation Q, no matter how arbitrary. Let Q, for example, be an equivalence relation, one of the equivalence classes of which is the set containing Heck's shoes, his daughter Isobel, the blackboard in Emerson 104, and perhaps some other things. Then define

$$\text{dud}(a) = \text{dud}(b) \text{ iff } aQb$$

Heck expresses doubt as to whether, for any relation Q, no matter how odd, there really are such things as duds, so defined (even assuming Frege's further condition to be fulfilled.) Insofar as the question, whether duds really exist, has any meaning, I myself see no reason to doubt that they do, unless some further implausible properties are claimed for them. The question, I would say, is whether duds are something that we have any reason to be concerned with. (The answer, as far as I can see, is that we do not, unless there is more to be said about them.) The apparent idleness of duds bears, however, on several further issues that are frequently discussed.

¹¹ Gottlob Frege, *The Foundations of Arithmetic*, sections 64-66.

The first is that views of the kind I am advocating are sometimes said to employ a “thin” or “minimal” sense of ‘exists.’ It is worth asking in what way this is so, and what the relevant “thinness” amounts to. One way in which this charge may seem to be correct is that I am claiming that the only general notion of existence, entailed by all existential assertions in any domain, is the purely formal notion captured fully by the logical rules governing the existential quantifier. This notion of existence may indeed seem “thin.” But before accepting it as a distinctive mark of the view I am proposing that it relies on a thin notion of existence, we should ask what kind of thicker notion it is being contrasted with.

The thicker or more robust alternative that is envisaged might be existence in the spatio-temporal world of objects, causes and effects. This is what is claimed when we say that mountains and elementary particles exist. I am of course not denying that this is the crucial content of these important existential claims. What I am saying is that the spatio-temporal world and the objects in it are not the only things that we can make true statements about, and that existential statements about other domains, such as statements about the existence of numbers, or rights, need not entail or presuppose claims that these things have spatial location or enter into causal relations.

A second way in which the view of existence I am advocating may seem to be “thin” or “minimal,” might be by contrast with a more robust metaphysical conception that might be thought to be entailed or presupposed by existential claims in any domain. I am denying that we need be concerned with such a conception of existence. But I want to resist the idea that this makes the conception of existence I am relying on implausibly “thin” or “minimal.”

Against the suggestion that this is so, I would note first that if, as I maintain, the only content common to existential claims in any domain is purely formal, and there are no metaphysical conditions of existence entailed by existential statements in all domains, this does not render statements about the existence of objects in the physical world “thin” or “minimal.” The “thickness” of the existential claims made by such statements is provided by the idea of that world itself. For physical objects to exist is for them to have spatio-temporal location, to have various physical properties, and interact with other objects in certain ways. That is to say, the relevant “thickness” is domain-specific. By the same token, then, the kind of thickness that is relevant to existential statements about numbers is provided by the structure of the relevant mathematical realm, not by some supposed metaphysical status that gives numbers a further degree of reality. For numbers to exist is for them to stand in various relations with sets, and with other numbers, to be the solution to equations and so on.

Even directions have a role to play in our talk both about abstract geometry and about lines, streets, and the movement of objects in the physical world. So even existential claims about directions have a degree of “thickness.” By contrast, the existence attributed to duds merely by the definitions I have mentioned seems thin because, unless more is said about them, there is nothing for their existence to come to: they do not do anything, or stand in any interesting relation to anything else. This is part of what is ensured by the meaninglessness of the relation Q .

On the basis of what has been said about them so far, duds do not belong to any domain. Where there is a domain, we might give objects introduced some role, by specifying how they enter into relations with other entities in that domain. If so, the claim

that they exist might have some “thickness.” But introducing these objects might still be pointless. Talking about them might bring no advantage from the point of view of our purposes in thinking about that domain. The following seems to be a plausible maxim in any domain: do not introduce new entities in a domain unless they have some significant role in that domain. Even if this maxim is valid for every domain, we would misconstrue the rationale for it if we understood it as a general principle of ontological economy. The point is not that it is better to have a universe that contains fewer entities rather than more. It is just a maxim of good sense. What is the point of talking about things that have no significant role in describing or explaining phenomena of the kind in question?

In the domain of natural science, the relevant role for entities is in the explanation of natural phenomena. So Harman’s explanatory requirement makes good sense: we have reason to be committed to the existence of things of a certain sort only if they play a role in explaining what happens in the natural world (including our experience of it.)¹² But this maxim is specific to the domain of natural science. It does not apply, as Harman’s explanatory requirement is often held to apply, to every domain, for example to the normative domain, or to mathematics. The relevant maxims for these domains will be, like the version of Harman’s requirement that I have just stated, domain specific. We have reason to introduce new terms denoting numbers (for example, imaginary and complex numbers) just in case these are useful in provide a more coherent and satisfactory account of the relevant parts of mathematics. And we have reason to

¹² Gilbert Harman, *The Nature of Morality*, Chapter 1.

introduce additional normative concepts and relations just in case these allow us to give a more coherent and satisfactory account of normative matters.¹³

Harman's explanatory requirement in the form in which he stated it (that we have reason to be ontologically committed only to things that serve in the best causal explanations of our experiences) applies to these other domains only if they are all taken to be about the same thing: the natural world. As I have said, I do not believe we should accept the idea that this world is the only thing we can speak sensibly, and truly about.

As I have also said, it is part of Quine's holism that mathematics and science are about the same thing: mathematics is just the most abstract part of our overall physical theory. Given this assumption, it makes sense for him to say that we have reason to accept ontological commitment to numbers because statements quantifying over numbers are needed in the formulation of the best scientific theories.¹⁴

¹³ In all these domains the explanatory task that justifies the introduction of new entities or concepts is the task of explaining the phenomena in question: what happens in the physical world, and the mathematical or normative facts. Explaining our reactions to, or beliefs about, these facts is a different matter. What I am suggesting as a condition for commitment to the existence of entities of a given type is similar to what Crispin Wright calls "the width of their cosmological role," by which he means "the extent to which citing the kinds of states of affairs with which it deals is potentially contributive to the explanation of things *other than*, or *other than via*, our being in attitudinal states which take such states of affairs as object." *Truth and Objectivity*, p. 196. The reference to "cosmological" role might suggest that the explanation in question must be causal explanation, of events in the natural world. But Wright goes on to say that this need not be so: "It is not my intention that that the Wide Cosmological Role constraint should be satisfiable only by *causally active* states of affairs, nor even that the explanations involved have to be causal. The overarching point, remember, is that there be a *wider range of intelligible and legitimate uses of the relevant state denoters than can be generated merely by the minimal truth aptitude of a discourse*. In principle, therefore, any additional kinds of context featuring the state denoters are significant, and interesting further distinctions may remain to be drawn depending on what the additional kinds of uses are." *Ibid.*, p. 198.

¹⁴ See Quine, ----; also Putnam, *Philosophy of Logic*, esp. Section VIII .

Seen from outside Quine's view, this seems strange. If the best scientific theories require commitment to particles or forces of a certain kind as elements of the physical world, then we have reason to be committed to such things. But the indispensability for physical science of quantification over numbers or sets, if true, is of a different order. Science requires the use of mathematics, but does not require the supposition that mathematical entities are part of the physical world. Quine's holism seems implausible insofar as it seems to have this implication.

The view I am advancing offers what seems to me a more plausible way of describing these matters. On my view, the existence of numbers and sets is properly settled entirely by the relevant mathematical criteria. Whether we have reason to quantify over imaginary numbers, for example, is determined by their role in the relevant part of mathematics. What the indispensability of mathematics, or some particular part of mathematics, for science does is not to give us reason to accept the existence of numbers, or of imaginary numbers. Rather, it gives us one kind of reason to be *concerned* with the domain of mathematics, or with that particular part of mathematics.

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 science, 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?¹⁵ 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 answer is that I see no such reason.

Let me turn now from these general remarks about ontology to particular questions raised by the idea of irreducibly normative truths. Contrary to what is sometimes said, belief in such truths does not involve commitment to any special entities.

¹⁵ There is also the question of whether the domains I have listed are in fact single, unified domains. Natural science has a certain unity because the various sciences are taken to be accounts of a single natural world. But is mathematics, for example, similarly unified? Do all the entities referred to in various parts of mathematics belong to single unified “world” of abstract objects? Or might there be distinct domains corresponding to various subfields of mathematics? This is a mathematical question, on which I take no position. I am grateful to Charles Parsons for calling this question to my attention.

The essential element in normative statements is not a term referring to an entity, but a relation: the relation 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 .¹⁶

This sounds like a three-place relation, but it contains an implicit universal quantification over agents: it holds that p is a reason for *any* agent in c to do a . And we want not only to make not only general claims of this kind but also to claim, of a particular agents x , that some fact p is a reason for him or her to do a in c .¹⁷ So the underlying relation must be a four-place one $R(p, x, c, a)$. I will, however, be concentrating on general claims, of the form $(x)R(p, x, c, a)$, which I will abbreviate $\mathbf{R}(p, c, a)$.

This emphasis on the generalized 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 must contain facts about the agent's desires in order for $\mathbf{R}(p, c, a)$ to be true. Also, this emphasis on general claims about reasons does not entail a commitment to the view that something is a reason for an agent only in virtue of certain facts about his or her situation,

¹⁶ The relational character of claims about reasons has been noted by Jonathan Dancy and by Terence Cuneo. See Dancy, *Ethics Without Principles* (Oxford: Oxford University Press, 2004) especially Chapter 3 pp. 38ff, and Cuneo, *The Normative Web* (Oxford: Oxford University Press, 2007), p. 65. Cuneo seems to have in mind a three-place relation (or perhaps a four-place one including a place for the agent.) Dancy mentions two two-place relations: the favoring/disfavoring relation and the enabling/disabling relation. A consideration can be a disabler if its presence undermines what otherwise would be a consideration favoring an action, but does not itself favor or disfavor that action. As will be clear below, and in Lecture 5, my three-place relation $\mathbf{R}(p, c, a)$ is intended to encompass both of these relations. Considerations c that are required in order for $\mathbf{R}(p, c, a)$ to hold are enablers in Dancy's sense.

¹⁷ I am indebted to Doug Lavin for very helpful discussion of this issue. I am not certain that he would be satisfied by my way of handling it.

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 the formulation I have given allows for the possibility that the identity of the agent may figure in the conditions c under which a given p is a reason for someone to do a .

The things that can occupy the first position in a statement, $\mathbf{R}(p, c, a)$ —the things that are reasons—may themselves be normative facts, but they need not be. Many, if not most, reasons are ordinary physical or psychological facts. 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 not the things that are reasons but the relation of being a reason for something. It thus not a matter of ontology *in Quine's sense* (the things quantified over), but rather of what Quine called “ideology” (the predicates employed.)¹⁸

This shift away from ontology (in Quine's sense) avoids one possible objection to irreducibly normative truths—a supposed commitment to strange entities. But it highlights two other potential difficulties having to do with the relation between normative facts and natural facts. Normative claims are not—as may be the case with claims about sets—simply about a special distinct realm. These claims, such as the claim I just mentioned about sharp objects, make claims about a relation among components or features of the natural world: facts (such as the sharpness of the metal), agents in certain situations, and actions these agents might perform.) The relation \mathbf{R} seems, therefore, to

¹⁸ This reflects Quine's view, which might be described as a form of nominalism about properties, that predicates in themselves have no ontological significance. Others may well see the use of the relation \mathbf{R} as raising an (ontological) question about what property corresponds to it, as I go on to say.

correspond to a (normative) relational property of these components and features of the world, and it may be wondered what this property is.¹⁹ Second, normative truths are distinct from (not entailed by) non-normative truths, but are linked to them in two significant ways: they vary when naturalistic facts vary, and they cannot vary as long as the naturalistic facts remain the same. These relations of covariance and supervenience need to be explained. To address these questions, it will be helpful to begin by looking more closely at the relation between the normative and the non-normative.

It is widely believed, by both cognitivists and non-cognitivists about the normative, that there is an important distinction, sometimes called an “unbridgeable gap,” between normative and non-normative judgments—between “facts” and “values,” as it is often put. But the idea that there is such a distinction has been challenged in a number of ways.

This idea is naturally expressed as the thesis that no normative statement is entailed by any set of non-normative statements, but in order for this to be the case the relevant class of normative statements needs to be specified more carefully.²⁰ Let F be a statement agreed to be factual, and N any statement agreed to be normative. What, then, about $F \vee N$? It follows logically from F. So if no normative judgment can follow from a factual one then it must be factual. But from $F \vee N$ and $\neg F$ one can deduce N. So $F \vee N$ cannot be factual if the thesis of non-derivability holds.

¹⁹ 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.” *Ethics: Inventing Right and Wrong*, p. 38.

²⁰ The following problem was noted by A. N. Prior. See his *Logic and the Basis of Ethics* (Cite??)

The existence of “thick” ethical concepts such as “cruel” is also often mentioned as showing that facts and values are inextricably intertwined.²¹ 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 conclusions, for example—often involves appeal to what are clearly values, such as simplicity, clarity and the like. So some paramount factual statements seem to depend on claims about value.²²

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

²¹ The term “thick ethical concepts” was introduced by Bernard Williams. See *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press, 1985) pp. 129-131. Hilary Putnam, mentions cruel as a counterexample to the idea that there is “an absolute fact/value dichotomy.” *Beyond the Fact/Value Dichotomy* (Cambridge, MA: Harvard University Press, 2002), p. 35. He does believe, however, that there is a *distinction* between the two (pp. 9-13.) Unlike a dichotomy, a distinction need not be omnipresent: it need not be the case that every item of the relevant kind falls on one side or the other.

²² See Putnam, *Beyond the Fact/Value Dichotomy*, pp. 135-145.

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.²³ 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 this requires some refinement.

I have said that the characteristic element in normative judgments is the relation $\mathbf{R}(p, c, a)$: “p is a reason for a anyone 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 from the non-normative claim that p is not the case that the normative claim $\mathbf{R}(p, c, a)$ is also false. This might seem to be a case of inferring a normative conclusion from purely non-normative premises. The essentially normative content of \mathbf{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 to do a. So I will understand \mathbf{R} in this subjunctive form and take what I will call a *pure normative claim* to be a claim of the form $\mathbf{R}(p, c, a)$, where \mathbf{R} is understood in this way.²⁴

²³ Putnam, *Beyond the Fact/Value Dichotomy*, pp. 19-24.

²⁴ I believe that such claims are what Kit Fine calls “world-bound normative conditionals” in “The Varieties of Realism,” pp. 272-273.

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 non-normative claims, it is a further claim that some pure normative sentence is true.²⁵

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. Mixed normative claims need not involve “thick concepts.” In everyday English, “p is a reason for x to do a” is a mixed claim, since cannot be true unless p is true. It also presupposes that there are conditions c such that the pure normative claim R(p, x, c, a) is true and x is in circumstances c. But the ubiquity of mixed normative claims is entirely compatible with the autonomy of the normative as I am formulating it.

The idea that there is a “logical gap” between normative and non-normative claims can seem puzzling, because 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?

²⁵ 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.

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

(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 the 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.

The distinction between normative and non-normative claims is most likely to seem like one that it is difficult to “bridge” if we focus on 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 often mentioned in this context, such as claims that something is good or that some action is morally wrong. These claims appear simply to assign to their subject some normative property, and the gap in question appears to be 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 non-normative properties. 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 normative/non-normative 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 other objections mentioned 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 indicate that the proposition in question is true. Beyond this, our reasons for accepting a scientific theory may, or may not, be based on further reasons that are not all truth-related.²⁶ 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 claims that we have sufficient reason for accepting a theory, or for counting a proposition as true, and claims made by that theory or proposition. The former may be normative even when the latter are not.²⁷

Nor does the intermingling of normative and non-normative elements in “thick ethical concepts” indicate that there is no distinction, in general, between the normative and the non-normative. 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 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.

²⁶ For discussion, see Joseph Raz, “Reasons: Practical and Adaptive.”

²⁷ A point made by Allan Gibbard. See *Wise Choices, Apt Feelings*, p. 34.

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 further 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 normative elements are central to the concept of cruelty. They give the content its point and explain 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.²⁸ 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.

Pure normative claims are not entailed by non-normative claims. But as I have said it is widely agreed that the normative is nonetheless linked with the non-normative in two significant ways. The first is covariance. Normative facts can depend on certain non-normative facts and vary when these non-normative facts vary. The second is supervenience. 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.²⁹

To understand the phenomena of covariance and supervenience it is important to be clear what kind of normative claims they involve. The normative facts that can vary as non-normative facts vary are facts that consist in the truth of mixed normative claims, such as the claim that someone has a reason to do a certain action, or that a particular consideration is such a reason. 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 were 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 mixed normative facts depend on non-normative ones, and which ones they depend on is a

²⁸ In Putnam's terms, it is a distinction, not a dichotomy.

²⁹ The problem is emphasized by Mackie, *Ethics: Inventing Right and Wrong*, p. 41; Blackburn, *Essays in Quasi-Realism*, pp. 114-148; and Frank Jackson, *From Metaphysics to Ethics*, pp. 118-125.

normative matter, specified by normative facts consisting in the truth of pure normative claims.³⁰ The truth of pure normative claims, by contrast, does not depend on, or co-vary with, non-normative facts.

Nor do pure normative facts vary “on their own,” without variation in non-normative facts. Given that they do not, the mixed normative facts which depend on them also supervene on non-normative facts. This again is a normative matter, a case of normative necessity.³¹ This seems evident from reflection on what pure normative truths are.³² But it does not seem to me, on reflection, to be something that we should find puzzling. It might seem puzzling if it were the very same normative facts that can vary when non-normative facts vary but cannot vary independently. But, as we have seen, this is not the case. It is mixed normative facts that co-vary with non-normative facts. Pure normative truths do not vary at all. Given that pure normative facts are not are not

³⁰ If this is correct, then it seems to follow that the dependence of the normative on the non-normative is quite different from the dependence of the mental on the physical. Normative judgments make claims about (the normative significance of) non-normative facts. That normative facts co-vary 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 “mentalist” truths do not specify which physical states mental states depend on and co-vary with. Confidence that the mental co-varies 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.

³¹ Kit Fine argues that the necessity of (some) normative claims is distinct from and not reducible to metaphysical necessity. See “The Varieties of Necessity,” Gendler, Gendler-Szabo and Hawthorne, eds., *Conceivability and Possibility* (Oxford: Oxford University Press, 2002), pp. 253-281. I would like to believe, although I am not certain, that his argument supports what I am maintaining here.

³² 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.

contingent in the most obvious way—that is, dependent on contingent facts about the natural world—why should we expect them to be contingent in some further sense?

I turn now to the question of normative properties. 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 it would be odd to think that there are normative properties—features “of the world”—corresponding to these concepts. 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 naturalistic.³³ 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 how the characterization of the property corresponding to a concept will go beyond what is specified by that concept itself. In some cases, having 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 other cases, however, 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.

³³ 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.

Consider first a concept of a natural kind, 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 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 concept is simply 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 appropriate guilt 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 or 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 the chemical composition of water. 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.³⁴ If this is right, then this is another case in which the minimal understanding of a property is insufficient. The further characterization of a natural kind such as water is scientific—a matter of chemistry and physics. The further 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 $R(p, x, c, a)$. 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 it would seem that no informative 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 having the chemical composition H_2O . What then should we say about the property signified by the concept, H_2O ? Perhaps this

³⁴ 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.

concept is fundamental: the property of being H₂O is just that, being H₂O. 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 earlier 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 in later lectures.

Allan Gibbard, although he accepts the idea that there are normative concepts, objects to the idea that there are normative properties, which he would find metaphysically odd.³⁵ He does not say exactly why, and it may be that he would not object to the idea of normative properties in the minimal sense I have proposed.³⁶

³⁵ *Thinking How to Live*, pp. 29-34.

³⁶ Simon Blackburn accepts the idea of normative properties in this sense. He writes, "There is no harm in saying that ethical predicates refer to properties, when such

Gibbard does not object to the idea of normative facts in a similarly minimal sense. He writes 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.”³⁷

Nonetheless, Gibbard proposes that the predicates corresponding to normative concepts are naturalistic. Discussing the normative concept *good*, he says that the property signified by this concept is the set of those things that are or would be good, contingencies aside. Transferring this to the relation of “being a reason for,” the suggestion would be that the property signified by this relation would be the quadruples $\langle p, x, c, a \rangle$ such that p would be a reason for a person, x , in situation c to hold attitude a if p were the case and x were 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

properties are merely the semantic shadows of the fact that they function as predicates.” “How to be an Ethical Anti-Realist,” in *Essays in Quasi-Realism*, p. 181.

³⁷ *Thinking How to Live*, pp. 182-183.

concept of a reason cannot be analyzed in non-normative terms.³⁸ 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 x in situation c to do a is for there to be some q such that x has a desire whose object is q and the truth of p is part of what explains how x's doing a promotes q.³⁹

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 is best understood in terms of reasons (as I would agree), then his account preserves normativity as long as it captures the idea of a reason.⁴⁰ One may think (as I do) that the naturalistic character of his account prevents it from doing this. But beyond merely asserting that his account cannot preserve the normative character of reasons, we should consider the grounds Schroeder offers 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.⁴¹ 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 to be the property of being a reason in part because it did nothing to explain the features that reasons have. Schroeder believes that his account does just this. Specifically, he

³⁸ *Slaves of the Passions*, p. 65.

³⁹ *Ibid.*, p. 59. I have modified Schroeder's definition slightly to fit my statement of R(p, x, 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.

⁴⁰ *Ibid.*, pp. 79ff.

⁴¹ 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.

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 has to say about covariance and 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."⁴² (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 their psychologies. But this claim 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

⁴² *Slaves of the Passions*, p. 2

that what explains the difference between Ronnie's reasons and Bradley's is the fact that Ronnie enjoys dancing and Bradley does not.⁴³ 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.⁴⁴ 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 for a person reason to perform given action must itself be part of that reason.⁴⁵ If this view were correct, then on a Humean theory a full statement of every reason for action

⁴³ 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.

⁴⁴ 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?

⁴⁵ *Slaves of the Passions*, p. 23.

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."⁴⁶ 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

⁴⁶ *Slaves of the Passions*, p. 27.

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."⁴⁷ One way to decide between these alternatives would be to imagine cases in which Ronnie would take pleasure in dancing

⁴⁷ *Slaves of the Passions*, p. 147.

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. 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.⁴⁸

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

⁴⁸ *Ibid.*

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.”⁴⁹ Suppose we know the set of quadruples $\langle p, x, c, a \rangle$ such that a language user assents to $R(p, x, c, a)$. 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 $R(p, x, c, a)$ holds. In order to know whether the language user assents to $R(p, x, c, a)$ 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 $R(p, x, c,$

⁴⁹ Jackson, *From Metaphysics to Ethics* (Oxford: Oxford University Press, 1998), p. 127.

a) as relevant to 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."⁵⁰ The property minimally signified by R, on my view, is not a "normative idler." To claim that R(p, x, c, a) holds is precisely to claim that (for x in c) p justifies doing 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. But, as I have explained, 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 natural or special metaphysical reality in order to have the significance that we commonly grant them.

Given the limited 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

⁵⁰ *Ibid.*

quasi-realism. Both Gibbard and Blackburn allow for, or even embrace, the idea of normative claims being true in what they see as 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.⁵¹

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.

⁵¹ See Gideon Rosen, "Blackburn's *Essays in Quasi-Realism*," *Nous* 32 (1998), pp. 386-405, and Jamie Dreier, "Meta-ethics and the Problem of Creeping Minimalism," *Philosophical Perspectives* 18 (2004), pp. 23-44. Dreier suggests (p. 37) that the difference between realism and expressivism lies in the fact that the former, but not the latter appeal to normative *properties* in order to explain certain phenomena. So perhaps he would place my view on the "expressivist" side of this distinction (as I would not.)

