

- FitzPatrick, W.J. (2014). Why there Is No Darwinian Dilemma for Ethical Realism. In *Challenges to Moral and Religious Belief: Disagreement and Evolution*, edited by Michael Bergmann and Patrick Kain. Oxford: Oxford University Press.
- FitzPatrick, W.J. (2015). Debunking Evolutionary Debunking of Ethical Realism. *Philosophical Studies* 172(4): 883–904.
- Greene, J. (2008). The Secret Joke of Kant's Soul. In *Moral Psychology*, vol. 3: *The Neuroscience of Morality*, edited by Walter Sinnott-Armstrong, pp. 35–80. Cambridge, MA: MIT Press.
- Griffiths, P. and Wilkins, J. (Forthcoming). When Do Evolutionary Explanations of Belief Debunk Belief? In *Darwin in the 21st Century: Nature, Humanity and God*, edited by Philip Sloan. Notre Dame: Notre Dame University Press.
- Haidt, J. (2001). The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment. *Psychological Review* 108(4): 814–834.
- Joyce, R. (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Joyce, R. (2013). Irrealism and the Genealogy of Morals. *Ratio* 26(4): 351–372.
- Kitcher, P. (2006). Biology and Ethics. In *The Oxford Handbook of Ethical Theory*, edited by David Copp. Oxford: Oxford University Press.
- Kitcher, P. (2011). *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Machery, E. and Mallon, R. (2010). Evolution of Morality. In *The Moral Psychology Handbook*, edited by John M. Doris, pp. 3–46. Oxford: Oxford University Press.
- Merritt, M., Doris, J., and Harman, G. (2010). Character. In *The Moral Psychology Handbook*, edited by John M. Doris, pp. 355–401. Oxford: Oxford University Press.
- Sabini, J. and Silver, M. (2005). Lack of Character? Situationism Critiqued. *Ethics* 115: 535–562.
- Shafer-Landau, R. (2012). Evolutionary Debunking, Moral Realism and Moral Knowledge. *Journal of Ethics and Social Philosophy* 7(1): 1–37.
- Sinnott-Armstrong, W. (ed.) (2008a). *Moral Psychology*, vol. 1: *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Sinnott-Armstrong, W. (ed.) (2008b). *Moral Psychology*, vol. 2: *The Cognitive Science of Morality*. Cambridge, MA: MIT Press.
- Sinnott-Armstrong, W. (ed.) (2008c). *Moral Psychology*, vol. 3: *The Neuroscience of Morality*. Cambridge, MA: MIT Press.
- Street, S. (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies* 127: 109–166.
- Street, S. (2008). Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About. *Philosophical Issues* 18: 207–228.

What's to be Said for Moral Non-Naturalism?

TERENCE D. CUNEO

In the not too distant past, many philosophers summarily dismissed moral non-naturalism. Though these philosophers acknowledged that non-naturalism has an impressive pedigree, having been defended by figures such as Plato and G.E. Moore, they charged it with being unmotivated, ontologically exorbitant, even “spooky” (see Jackson 2012). Christine Korsgaard, for example, describes the view as one according to which there are moral facts “wafting by” that we somehow grasp (1996, 44). Allan Gibbard characterizes it as requiring of us special powers to “peer into a special realm of normative fact” (1990, 284). Though these images of non-natural facts as belonging to some ethereal domain continue to operate on the thinking of moral philosophers, they no longer exert the influence that they once did. Thanks to the efforts of philosophers such as Russ Shafer-Landau (2003), David Enoch (2011), and Derek Parfit (2011), non-naturalism is once again a view that many philosophers take very seriously.

My project in this chapter is to give you a feel for why an increasing number of philosophers take non-naturalism seriously, presenting a two-stage argument in its favor. The first part of the argument presses the point that moral naturalists must defend some form of reductionism according to which moral properties are “nothing over and above” natural or descriptive properties. Vindicating reductionism, I contend, is considerably more challenging than many have assumed. The second part of the argument develops the claim that once we specify what moral facts would have to be like, there is reason to hold that some moral facts are as non-naturalists believe.

At the end of this chapter, I will indicate why I believe that the argument I'll present is importantly incomplete. Before moving forward, however, I should flag two other respects in which the argument has limited ambitions. First, any full defense of non-naturalism would have to specify how non-naturalists can respond to important challenges to their view, of which there are several. While I believe that such a defense can be provided, I'll not attempt to offer it. Moreover, any defense of “robust” non-naturalism would have to explain why thinner versions of the view, which claim that while there are non-natural moral truths,

these truths have “no positive ontological implications,” are unsatisfactory (see Parfit 2011, 478; see also Scanlon 2014). I’ll not attempt to do that on this occasion. Instead, I’ll assume that non-naturalism makes substantive and controversial ontological commitments, specifically concerning the nature of moral facts.

What is Moral Non-Naturalism?

Let us begin by identifying several points on which moral naturalists and non-naturalists agree. First, philosophers of both sorts maintain that there are moral facts, such as *that it is wrong to cheat others simply for gain*. Second, both (as I shall think of them) are moral realists, agreeing that some moral facts are objective, in the sense that they hold independently of what we value or what conventions hold. According to realists, the fact to which I just adverted does not exist because some of us are against cheating or because there are social conventions that prohibit cheating for gain. Finally, both naturalists and non-naturalists agree that moral facts do not float free of the natural realm but bear intimate relations to it; actions are wrong, these philosophers agree, because of the way things are at the natural or descriptive level.

Though naturalists and non-naturalists agree on this much, they disagree about the nature of moral facts. Moral naturalists hold that moral facts are nothing over and above natural ones. Moral non-naturalists demur, holding that some moral facts depend on but are not themselves natural facts. While many philosophers believe that this difference between naturalism and non-naturalism is deep and important, offering an informative specification of what distinguishes natural from non-natural facts has proven extraordinarily difficult. All the proposals in the literature are problematic.¹

Given our topic, it would be helpful to have an informative account of what distinguishes natural from non-natural facts. While I would very much like to have such an account, I have none to offer. In its absence, I propose to employ a less-than-fully-informative but nonetheless expedient route, which is to assemble two lists. The first includes terms that putatively refer to entities that most parties to this debate would identify as being natural, perhaps because they are such as to play explanatory roles in the natural sciences or are reducible to such features. The second includes putatively referring terms that don’t make it on to the first list. According to this strategy, if a term doesn’t make it on to the first list, then it is a good candidate for being an expression that might refer to something non-natural.

Nothing about this approach, I should emphasize, implies that terms on either list refer. Nor does this approach imply that if these terms do refer, those on the first list refer to natural entities, while those on the second list refer to non-natural ones. After all, for all I say here, Berkeleyan idealism (which holds that everything that exists is nonphysical) might be true, in which case many of the terms on the first list would either not refer or would refer to non-natural entities. Or, alternatively, reductive physicalism might be true, in which case many of the terms on the second list would not refer or would refer to natural entities. The role of compiling these lists, then, is heuristic; it is simply to furnish promising candidates for entities that would be natural or non-natural entities (if there are any).

¹ Copp (2007, ch. 1), Cuneo (2007), McPherson (2015), and Shafer-Landau (2003) offer different ways of understanding the distinction.

Let’s consider what sorts of term these lists might include. A version of the first list should probably include terms that putatively refer to objects of various sorts, such as the planet Mars, quarks, and elephants. Likewise, it should probably include terms that refer to properties such as *having negative charge*, *being in pain*, and *having the atomic number 8*. And, finally, it should probably include terms that refer to concepts that correspond to these properties, such as *HAVING NEGATIVE CHARGE*, *BEING IN PAIN*, and *HAVING THE ATOMIC NUMBER 8*. The second list might include terms that are not included on the first, such as expressions that putatively refer to objects such as God. This list might also include terms that refer to properties such as *being sublime*, *being a promise*, or *counting in favor of*. And, finally, it should include terms that refer to concepts that correspond to these properties, such as *BEING SUBLIME*, *BEING A PROMISE*, and *COUNTING IN FAVOR OF*.²

Non-naturalists hold that terms that refer to moral concepts and moral properties belong on the second list. The rationale that non-naturalists offer for this claim is that moral concepts and properties are importantly different from the sorts of concept and property that belong on the first list. Think, for example, of moral concepts such as *BEING WRONG* and *BEING UNJUST*. These concepts are not such as to play any interesting explanatory roles in our present sciences, including physics, chemistry, and biology. Nor do we have any reason to believe that they would be fit to play any such role in a perfected or idealized version of these sciences, since we have no clear idea what these idealized sciences would be like. Nor, moreover, do these concepts appear to be reducible to concepts that are such as to play explanatory roles in the sciences, such as *HAVING NEGATIVE CHARGE*, *HAVING THE ATOMIC NUMBER 8*, or *BEING FITNESS ENHANCING*. As philosophers since G.E. Moore have emphasized, attempts to argue that a moral concept such as *RIGHT* is identical with a natural concept such as *BEING SUCH AS TO ENHANCE EVOLUTIONARY FITNESS* have come to grief, since there is nothing incoherent or inapt in thinking that an action enhances evolutionary fitness but is not right. (Selective cheating, for example, might enhance human evolutionary fitness but is, in the ordinary case, not right.)

The line of argument articulated in the last paragraph suggests that there is a more general point to be pressed against any attempt to identify moral with natural concepts. For, by all appearances, natural concepts have a certain function or job description, which is (roughly) to figure in theories that predict and explain occurrences in the natural world, such as happenings on the quantum level or the behavior of gasses. Moral concepts, however, appear to have a very different job description. Their primary function is not to figure in theories that predict and explain what happens in the natural world, but rather is *practical*: they are employed to help guide decisions and evaluate behavior. For example, by noting that you have a decisive moral reason to return a book that you’ve borrowed from me, you can determine how you ought to act. If you fail to do what you should, we can employ moral concepts such as *WRONG* or *INCONSIDERATE* to evaluate what you’ve done (or not done). On the basis of the application of these concepts, we can then determine whether to hold you accountable, admonish you, blame you, or the like for failing to return the book. When we do so, we are not trying to explain or predict your actions; we are evaluating them. We are thinking “ought thoughts.”

There is, then, a very general reason to believe that natural and moral concepts are fundamentally different, as they appear to play different roles in our lives. Since this claim

² Like Gibbard (2003; 2012), I hold that concepts can be non-natural. For a brief defense of this claim, see Cuneo and Shafer-Landau (2014).

is important to the following argument, it will be helpful to have a name for it. I'll call it the *explanatory/normative divide*. Having identified this claim, it is now important to highlight a distinction that was largely implicit in the preceding paragraphs – a distinction between *concepts*, on the one hand, and *properties* and *facts*, on the other.

To get a handle on this distinction, begin by noting that we often operate with a rough and ready distinction between thoughts and world – between what we think and what there is. Concepts belong to the “thought” side of this distinction, since (as I shall understand them) they are abstract and sharable ways of thinking about objects and properties. By employing concepts such as HORSE, NEIGHBOR, and SKITTISH, for example, I can form the thought THAT MY NEIGHBOR'S HORSE IS SKITTISH – this thought or proposition having as its constitutive components the concepts HORSE, NEIGHBOR, and BEING SKITTISH. Like many other thoughts, this thought purports to be about the world: in this case, my neighbor's horse. Once we see this, we can also see that concepts are not merely the subcomponents of propositions, they are also *referential devices*. They are that which enables thinkers to think about or refer to things such as horses and skittishness – if any such things there be.

Unlike concepts, properties belong to the “world” side of the thought–world distinction, since they are features that entities can have or exemplify. Suppose, to stay with our example, that my neighbor's horse actually is skittish. If it is, then this horse has the property of *being skittish*. When it does, we can say that it is a fact that *my neighbor's horse is skittish*, as many facts are simply situations or states of affairs in which things have properties. Understood in this way, facts often make propositions true. My thought THAT MY NEIGHBOR'S HORSE IS SKITTISH is made true by the fact that *my neighbor's horse is skittish*. If I also form the moral thought THAT MY NEIGHBOR HAS TREATED HIS SKITTISH HORSE WRONGLY, then this thought is true if and only if and because it is a fact that *my neighbor has treated his skittish horse wrongly*.

The reason that I am calling attention to this distinction between thought and world is this: often we think of the same thing by using different concepts. If you are baking a cake, you might have the thought that you need to turn up the heat in the oven that you're using. If, by contrast, you are writing a paper in physics, you might have the thought that scientists have failed to understand certain properties of mean kinetic energy. While in both cases you would be thinking different thoughts, you would nonetheless be thinking about the same thing – namely, mean kinetic energy – but in different ways. Otherwise put, while the concepts HEAT and MEAN MOLECULAR KINETIC ENERGY are different, they refer to the same thing. This observation opens the conceptual space for the following possibility, which is important in the present discussion.

The possibility is one according to which non-naturalists are correct to claim that moral concepts such as WRONG cannot be identified with natural or descriptive concepts but incorrect to hold that these concepts characterize non-natural properties or facts. This possibility holds not because moral concepts fail to refer at all, as error theorists believe. Rather, it is because it seems possible that moral concepts characterize only natural or descriptive properties or facts. To those familiar with debates in the philosophy of mind, the structure of this view should seem familiar, as it would be analogous to a view according to which concepts such as PAIN OR TASTING CHOCOLATY are not identical with concepts such as BEING IN SUCH-AND-SUCH BRAIN STATE, but nonetheless refer (only) to brain states. We might, of course, be completely ignorant of the fact that moral and mental concepts work that way. However, successfully referring to something by the use of a concept needn't imply that those who employ such a concept have correct views about how that concept works or to what it refers.

The debate between non-naturalists and naturalists has various dimensions, and in our discussion so far we have only touched upon some of them. Even if we do not explore the various dimensions of this debate – such as what motivates philosophers to accept one or another of these positions – we are in a position to see that a key point of disagreement will be whether it's plausible to believe that moral concepts, which by all appearances are not natural, refer to natural properties or facts. Non-naturalists must offer reasons for believing that they do not and naturalists must furnish reasons why they do. Who should we believe? In the next section, I discuss two naturalist positions, each of which employs different strategies to contend that, while moral concepts are non-natural, moral properties and facts are natural. The upshot of the argument I present is that these strategies are deeply problematic.

Two Reductionist Strategies

Suppose you believe that everything that exists is part of the natural world. If you also believe that there are moral facts, as moral naturalists do, then you are committed to the claim that they too must be part of the natural world, since they could be nothing over and above natural facts. In claiming that moral features are nothing over and above the natural world, moral naturalists endeavor to reduce moral features to natural ones.

As it happens, naturalists have not agreed about what it is for something “not to be over and above” the natural realm. According to some naturalists, such as Frank Jackson, the “not over and above” relation is simply that of identity (see Jackson 1998). Modified somewhat, Jackson's argument for the thesis that moral features are identical with natural ones runs as follows.³ Moral features such as *being wrong* do not float free of natural properties such as *causing pain*; rather, they supervene on them, as there cannot be a change in a thing's moral features without a change in its natural features. Now, suppose we were to bundle together all the natural features on which a given moral feature – say, wrongness – could supervene, creating a huge disjunctive natural property. Call this property N. N would be necessarily equivalent with wrongness, for, necessarily, something would be wrong if and only if it were N. But, says Jackson, we should accept the sparse theory of properties according to which properties that are metaphysically equivalent are identical. It follows from the sparse theory that wrongness is identical with N, which is a natural property. Since there is nothing unique about wrongness in this respect, it follows that every moral property is identical with a natural one.

When understood as a view about facts, Jackson's argument faces an especially pressing worry, which is its move from:

Fact A and fact B are metaphysically equivalent: necessarily, A exists if and only if B does.

to:

Fact A and fact B are identical.

³ I am omitting various details of Jackson's position, including the fact that he is an analytic naturalist, holding that moral and descriptive terms share the same semantic content.

One concern about this move is that it yields bizarre results. After all, if it were correct, it would follow that the fact *that there is such a person as God or it is false that there is such a person as God* is identical with the fact *that $2 + 3 = 5$* , since they are metaphysically equivalent. However, these facts certainly do not seem to be identical; one concerns numbers, the other doesn't (see Plantinga 2010).

Another worry, perhaps more pressing than the first, is that properties and facts can bear different sorts of relation to one another, including that of *being more fundamental* than another or *being such as to determine* another, even when they are metaphysically equivalent. However, Jackson's view, which takes metaphysical equivalence to be sufficient for identity, is insensitive to determination relations that might hold when one property or fact is more fundamental than another.

To illustrate, consider the fact *that sugar is sweet*. This fact is metaphysically equivalent to the fact *that an omniscient agent would believe that sugar is sweet*. But, on the face of things, this provides little reason to believe that these facts are identical. To the contrary, it seems that the second fact holds in virtue of the first; an omniscient observer would believe that sugar is sweet because it is. When applied to the moral domain, we can raise the same concern. Suppose, for argument's sake, that the properties *being morally obligatory* and *being such as to maximize happiness* are metaphysically equivalent. One possibility is that the latter property determines the former; another possibility is that they are identical. Short of denying that there are any determination relations at all, however, the fact that these properties are metaphysically equivalent would not itself imply that they are identical (any more than the fact *that sugar is sweet* is metaphysically equivalent with the fact *that an omniscient being would believe that sugar is sweet* establishes that these facts are identical).

For those who believe that facts or properties can determine one another, then, there are good reasons to be uneasy with Jackson's reductionist proposal. There is, however, another variety of naturalism that rejects the idea that the "nothing over and above" relation is best characterized in terms of identity. Philosophers who espouse this view maintain that to reduce moral features to natural ones is to show that natural features *constitutively explain* moral ones. To get a feel for this view, consider an example from the non-moral realm: the property *being a triangle* is constituted by *having, three, angles*. Something is a triangle because it has these constituents, but these constituents are not what they are in virtue of constituting a triangle. If this is right, reducing a property F to a property G is not so much a matter of establishing that F and G are identical as establishing that the constituents of G constitutively explain F.

Mark Schroeder, who defends this position, maintains that any putative reduction of this (or indeed any) sort will be successful only when a pair of conditions is met (see Schroeder 2007). For one thing, any object or property that is a candidate for being reduced will have "fixed features": features that we have strong reason to believe must be true of it. Most believe, for example, that to be in a state of pain is to feel a certain way (call this feel being "ouchy"). A successful reduction of pain states to brain states, then, must imply that pain states are "ouchy," this being a fixed feature of pain states. Moreover, any successful reduction must leave enough of our ordinary background picture of the world intact. A reduction is not successful if it implies that pain states are "ouchy" but that "ouchiness" is not a feel, but simply a tendency to wince or to curse. This broadly behaviorist construal of ouchiness is unacceptable because it is a deeply entrenched part of our background picture of the world that we can experience feels of various sorts. Could any putative reduction of the moral to the natural satisfy these two conditions? Philosophers such as Schroeder believe so. If we

understand reduction in terms of constitutive explanations, these philosophers hold that fundamental normative features, such as *being a reason*, might well reduce to natural features, such as *being desired*.

Let's consider the constitutive explanation strategy in more detail, focusing on Schroeder's proposal, since it is the most developed. Schroeder's case for his view begins with a simple thought experiment: imagine that there is going to be a party tonight, at which there will be dancing. Imagine, further, that while Ronnie loves to dance, Bradley hates it. On the face of things, the fact that there will be dancing at the party is a reason for Ronnie to go, since he loves to dance. In contrast, the fact that there will be dancing at the party is a reason for Bradley not to go, since he abhors it. If this is right, we know that in some cases, whether an agent has a reason to act in some way is determined by features of his psychology; specifically, his desires. Might all of an agent's reasons be explained in this way? Schroeder suggests that the answer is yes: for any fact whatsoever, if that fact is a reason for an agent to act, then we can explain it in terms of that agent's desires.

Put more formally, Schroeder's proposal is what he calls:

Reason For [a proposition] R to be a reason for X to do A is for there to be some [proposition] p such that X has a desire whose object is p, and the truth of R is part of what explains why X's doing A promotes p. (2007, 59)⁴

Stated more simply, **Reason** tells us that an agent's reasons to act are those propositions whose truth helps explain why acting in some way increases the likelihood of that agent's satisfying one or another of his desires.

Before critically engaging Schroeder's proposal, let me try to clarify it further. First, Schroeder's view is not that desires are reasons but that desires constitutively explain why certain features of the world – such as that there will be dancing at a party – are reasons. Thus understood, desires are "reason-explainers." Second, Schroeder has a very liberal understanding of the promotion relation: anything that raises the probability of an agent's satisfying a desire "promotes" the satisfaction of that desire (2007, 113). Third, an implication of this last point is that, for just about any action whatsoever, there will be a reason for an agent to perform it, since that agent will have *some* desire such that performing that action will raise the likelihood of his satisfying that desire. It does not follow from this last point, of course, that if you have a reason to perform some action, then you *should*, since reasons can have different weights. As stated, **Reason** tells us nothing about how much weight a particular reason might have for an agent (see Schroeder 2007, ch. 7).

It is easy to see that **Reason** is a very controversial proposal, since it is very inclusive. It implies, among other things, that you now have reason to chew on your car – the explanation being that since you desire to be healthy and chewing on your car will introduce healthful iron into your diet, the likelihood that this desire will be satisfied will be increased if you chew on your car. **Reason** also implies that, given that Ronnie has certain desires – such as not wanting to injure himself – he has reasons not to dance, and that, given some of Bradley's desires – such as wanting to be more physically fit – he has reasons to dance. This is odd, because the case of Ronnie and Bradley was supposed to be one that illustrates that

⁴ According to this formulation, reasons are propositions. Elsewhere, Schroeder speaks as if they are facts or events. While I believe that propositions, facts, and events belong to different ontological categories, when I sketch Schroeder's view, I'll not pay attention to these differences.

they have reasons to act differently. Given Schroeder's view, it does not illustrate this at all, since the fact that there is dancing at the party is, for each of them, both a reason to attend and a reason not to attend (see McPherson 2012).

Suppose, though, that we waive such initial concerns, assuming that **Reason** is correct. If it is, it will follow that an agent's having a reason to act in a certain way is reducible to natural features, namely, its being such as to raise the likelihood of that agent's satisfying a desire. The question is whether this proposal is satisfactory.

Let me offer two reasons for believing not. First, it is one thing to show that a particular kind of normative feature, namely, being a reason, is reducible to natural features; it is another thing to establish that *all* normative features are reducible to descriptive features, which is what a naturalist position must establish. Schroeder expresses optimism that if reasons are reducible to natural features, then normative features as such are reducible to natural ones, since we can constitutively explain all normative features in terms of reasons (2007, ch. 4). While I am not prepared to say that such a reduction is impossible, I see no reason for being optimistic about its prospects.

Here is one reason why: Suppose you have a moral right against me that I not attack you (and, hence, I have an obligation not to attack you). And suppose that I have the same right against you (and, hence, you have the same obligation toward me). A "reasons-first" approach of the sort that Schroeder favors must constitutively explain our having this right in terms of reasons. Moreover, on the assumption that we each have this right in virtue of some shared characteristic, it follows that, under Schroeder's position, there must be some desire such that each of us has it, and the likelihood of its being satisfied is increased by my not attacking you and your not attacking me. But it is (to put it mildly) highly contentious that there is such a desire. The desire to avoid bodily harm is not such a desire, for it is possible that I lack this desire; if that is the case, however, it doesn't follow that it is morally permissible for you to attack me.

The second concern about Schroeder's proposal is whether it satisfies its own strictures for being successfully reductive. Recall that, according to Schroeder, any putative reduction will be successful only when the following is true. First, any object or property that is a candidate for being reduced will have "fixed features": features that we have strong reason to believe must be true of it. Moreover, any successful reduction must leave enough of our ordinary background picture of the world intact. The attempt to constitutively explain triangularity in terms of *having, three, angles* seems to satisfy these criteria; it does no violence to our ordinary understanding of triangularity – to the contrary, it illuminates it. The attempt to constitutively explain the "ouchiness" of pain in terms of behavioral dispositions, in contrast, appears to fail these criteria, since pain appears to be not a behavioral disposition but a feeling. This attempt at reduction does great violence to our ordinary understanding of pain – such violence that an eliminativist about the mental (someone who denies that there are any mental events as we ordinarily think of them) could probably accept it. The question to ask is whether **Reason** more nearly resembles the mathematician's attempt to constitutively explain triangularity or the behaviorist's attempt to constitutively explain pain.

Reason, I believe, is closer to the behaviorist's attempt to constitutively explain pain, and for a very fundamental reason: **Reason** reduces reasons to an explanatory relation, since it tells us that a reason for you is simply what raises the probability that a desire of yours is satisfied (see Olson forthcoming). But normative relations such as *being such as to favor* and *justifying* are, by all appearances, not merely explanatory relations. After all, it seems possible for a fact – such as *that there will be dancing at the party tonight* – to help explain

another fact about me and my mental life but not to favor or justify my behaving in any way at all. For example, suppose the fact *that there will be dancing at the party tonight* raises the likelihood that some *belief* of mine is true. That hardly seems to count in favor of or to justify (to any degree) my going to the party, let alone to be what such a reason consists in. Contrast this with the case in which the fact *that there will be dancing at the party tonight* raises the chances that I'll have a *desire* satisfied. So far as I can see, this last proposal seems no better situated to constitutively explain reasons than the first. Neither seems to capture what it is for a fact to recommend, favor, or justify my going to the party. Otherwise put, **Reason** does not bridge what I earlier called the *practical/explanatory divide* (at least when this divide pertains not just to concepts but also to properties).

The Structure of Normative Facts

My project in the preceding section was to raise some general concerns about attempts to reduce moral features to descriptive ones. Jackson's strategy, we saw, not only yields bizarre results but is also insensitive to determination relations that may hold between properties or facts. Schroeder's proposal, in contrast, gives us no particular reason to believe that specifically moral features can be reduced to descriptive ones, and also appears not to satisfy the criteria for a successful reduction, being too revisionary in character. To point to the liabilities of these reductive proposals is, of course, not to establish that no reductive proposal could succeed. But the failure of these views, I believe, provides grounds for pessimism regarding future proposals.

In this section, I want to move from critiquing attempts by naturalists to reduce the moral to the natural to presenting a positive case for non-naturalism. As I indicated at the outset of our discussion, this case hinges on observations about the structure of normative facts: what moral facts would have to be like if they were to exist.

My point of entry into this topic is a recent discussion by William FitzPatrick (2008), in which FitzPatrick develops his own version of non-naturalist realism. He observes that naturalists such as Jackson have often developed their views by attempting to identify moral features with massive disjunctions of natural properties. I have pressed the point that such views are insensitive to determination relations that might hold between properties or facts. FitzPatrick takes the argument a step further, specifying the types of determination relation that hold when something is a moral fact.

Like Schroeder, FitzPatrick begins with a thought experiment. Consider, says FitzPatrick, an ordinary artifact such as the computer on your desk. Now ask yourself why that computer is a good one. Your answer will appeal to ordinary natural properties, such as its processing speed, its propensity not to crash, and so forth. For ease of reference, call these features of your computer XYZ. Though your computer is good in virtue of having these properties, its being good does not consist solely in its having them, for any number of things could have these properties without being good. Moreover, when it comes to artifacts, at least, good-making characteristics are relative to kinds; sharpness is a good-making quality in a knife, for example, but not in a coffee cup. FitzPatrick maintains that this point is important

not because...[it] translates directly to the ethical case – it does not, as at least some ethically good- or bad-making properties may plausibly have that status invariantly...but because it directs us to the more complex structure of facts about goodness, which does have an analog in

ethics. In the present case, the point is that the fact that this computer is good consists not only in the fact that it has XYZ, but in this *together with* the fact that XYZ is such as to satisfy the *standards of goodness S* for computers. (2008, 186)

By noting that the goodness of an artifact of some kind is determined by the standards of goodness for members of that kind, FitzPatrick does not wish to defend non-naturalism about such standards. Rather, he wishes to direct our attention to, as he puts it, the structure of goodness-involving facts, or as I would prefer to put things, the structure of normative facts.

FitzPatrick's argument appeals to parallel cases: To account for the normative properties of artifacts, we must appeal to normative standards, such as the standards of goodness for computers or knives, which explain how their natural properties contribute to their goodness or badness *qua* artifact. Likewise, to account for the moral properties of actions, we need to appeal to moral standards, which explain how an action's natural properties contribute to its having such properties as being morally permissible or impermissible.

Let's develop FitzPatrick's point. Imagine being present at a social gathering in which a senior colleague, who has a strong sadistic streak, publically humiliates a junior colleague with a clever but vicious put-down. Your senior colleague's behavior is morally wrong. But it is wrong not simply because it causes humiliation and social tension, which are ordinary natural states or properties, but also because there is the moral standard that:

It is wrong to humiliate others simply for pleasure.

I would add that this last point about the structure of moral facts extends to other cases, such as those involving etiquette and prudence. To see this, imagine that the same senior colleague who has acted badly has also acted rudely and irrationally. If his action is rude, however, then it is in virtue of his having violated standards of politeness that apply to his behavior. And if his action is practically irrational – say, because it deeply alienates those whose company he generally enjoys – then it is in virtue of his having violated applicable standards of prudence. This suggests that the structure of normative facts to which FitzPatrick draws our attention is pervasive; it is found not simply in the moral domain but also in the normative domains quite generally (cf. 2008, 188).

In sum, FitzPatrick defends three claims:

- (1) There are particulars, such as act-tokens and individual people, that have moral properties. A particular's having a moral property is a particular moral fact.
- (2) When a particular has a moral property, it has that property in virtue of there being some moral standard which that particular satisfies or fails to satisfy.
- (3) Moral standards are general moral facts.

While the first claim is accepted by nearly all moral realists, the other two are more contentious. For that reason, I'll revisit them in a moment. Before I do, however, it's worth asking this: Why would a view according to which there are moral standards that determine particular moral facts favor non-naturalism in any regard?

Perhaps the best way to see why is to work with a pair of contrasting cases. Consider a game such as ice hockey. Like most games, hockey is a normed activity, since it's governed by rules or norms for how to play it. For example, the rules or standards that govern the game tell us that a standard game is divided into three periods, at most only six players on

each team can play at once, a player is not allowed to enter the opponents' zone before the puck does, the puck is considered "dead" during play if it leaves the perimeter of the rink, and so on. Along with these standards are rights, responsibilities, and obligations that attach to participants in the game, namely, players, coaches, and officials. Players, for example, have a right to take a shot on goal but they do not have the right to determine whether such a shot counts as a goal (officials do); officials have the right to determine whether a player is off-side but coaches do not; officials have the responsibility of tracking how many penalties a player might have but players do not; and so on. These rights, responsibilities, and obligations are enforced in various ways. If a player, for example, fails to comply with them, he or she may be penalized, assessed a fee, or suspended. If this is right, the standards that govern the game of hockey play at least two roles: they determine not only whether someone is playing the game correctly or incorrectly but also whether an activity counts as hockey at all. For were some activity that had the trappings of hockey – a game played on a rink with sticks, a puck, and so forth – not to sufficiently conform to the standards of the game, this activity would not count as a game of hockey.

To my knowledge, no philosopher has advocated non-naturalism regarding the standards that govern hockey and the correlative rights, responsibilities, and obligations that attach to participants in the game. Still, it will be worthwhile to think about the ontological status of the standards themselves, if only to understand better why these standards have been assumed to be naturalistically acceptable. We can do this by asking whether the standards of hockey bear various "marks of the natural."

According to many, the paradigmatic mark of the natural is to be fitted to play an explanatory role in the sciences. This is why quarks, DNA, and nitrous oxide are rightly deemed natural: they bear this mark. The standards of hockey, however, certainly don't bear this mark, since no science would advert to these standards to explain anything in the natural world. Another mark of the natural to which some philosophers appeal when defending the naturalistic credentials of some entity is being causally efficacious. Once again, though, the standards of hockey don't seem to bear this mark. In fact, if such standards exist, they look, for all the world, to be not particular concreta that are causally efficacious, but abstract entities that exist but are not spatially located (though they might have a history if we brought them into existence). If these standards enter into causal explanations at all, they will do so only by being the contents of our propositional attitudes.

So far, the standards of hockey don't seem to bear the marks of the natural. But notice this: these standards have the property of being deeply contingent. By this I mean that not only is the existence of the game of hockey contingent, since we can easily imagine never having invented it, but also the game itself could (at least to some degree) be governed by different standards. We could imagine, for example, the National Hockey League modifying the game in such a way that during certain periods of the game, seven players from a team are allowed to play at once.

There is an explanation for the deep contingency of the standards of hockey: they owe their existence or implementation to our activity. It is we who implemented these standards, by performing acts of authorization. For, at some point, authorized parties performed speech acts in which they pronounced that hockey games would be composed of three periods, that no more than six players from one team would be allowed to play at the same time, and so on. When they pronounced these things, they implemented the standards that govern the game. (I don't mean to suggest that this is the only way to implement standards. It is probably one among several.) These standards, moreover, continue to be in force.

That is because we continue to care about them and enforce penalties when they are violated. If we were to cease to care about or enforce the standards of hockey, these standards would probably still exist, but they would fail to govern anyone's activities.

In making these observations, we are still a considerable distance from adequately explaining how the standards of hockey are or could be constitutively explained by paradigmatic natural features. In fact, I believe that there are important unresolved puzzles regarding how these standards could be so explained (see Cuneo 2014). My purpose in discussing the standards of games such as hockey, however, is not to build a case for holding that we should be non-naturalists about them; rather, it is to compare them to moral standards, which non-naturalists believe we do have reason to hold are not constitutively explained by paradigmatic natural features.

To that end, let us assume that the concerns raised earlier about views such as Jackson and Schroder's are weighty enough to place these broadly reductionist positions to the side. If we do, then we need to ask whether moral standards are plausibly viewed as being descriptive, bearing the marks of the natural. We've already noted that, according to many advocates of naturalism, the paradigmatic mark of the natural is being such as to play an explanatory role in the natural sciences. Are moral standards fitted to play such a role? On the face of things, no; moral standards seem no better situated to play such a role than the standards of hockey, since our best theories in biology, physics, chemistry, and the like do not advert to properties such as *being wrong* or *having a moral right*.

Naturalists such as Nicholas Sturgeon suggest, however, that even if moral facts are not such as to play explanatory roles in the natural sciences, they are nonetheless plausibly viewed as being causally efficacious (Sturgeon 1988). For example, if we wished to address the question why your senior colleague acted badly, we might advert to moral facts such as his being morally callous. In my view, Sturgeon might be correct to say that some moral facts, such as *that your colleague is morally callous*, could causally explain happenings in the world (see Cuneo 2006; see also Shafer-Landau 2003; Wedgwood 2007). Still, it would not follow from this that the moral standards themselves causally explain anything. Like the standards of hockey, moral standards look for all the world not like concreta, such as quarks or concatenations of chemicals, but like abstracta, which are neither spatially nor temporally located. As such, if they enter into the causal nexus, it is because we represent them in certain ways.⁵

It is at this point, however, that we come to a crucial difference between the standards of hockey and the standards of morality. Unlike the standards of hockey, many of the standards of morality do not look contingent at all. Instead, they look to hold necessarily, at least for beings like us in a world such as ours. If we hold fixed our present constitution and environment, it looks as if the following standards *must* hold:⁶ It is wrong to humiliate others simply for pleasure; is wrong to kill someone simply because she has inconvenienced you, and; is wrong to engage in recreational slaughter of fellow persons. One explanation of why they must hold is that it belongs to the essence of wrongness that killing someone simply because she has inconvenienced you is wrong. If such a view were correct, it would have important implications for how we could know moral standards, for we could know these standards by grasping important features of the essence of properties such as *being wrong*. If this were so, however, then our knowledge of the standards would not be via

5 A full vindication of this claim would, however, have to engage with Railton (1986), who claims that normative standards can enter into nomic explanations even when not the content of our beliefs.

6 This is compatible with these facts having *ceteris paribus* clauses built into them. More on this in a moment.

wholly empirical means or through the usual methods of the sciences, such as drawing generalizations by inductive means from particular cases. It would involve something like *a priori* insight into the way things are.

In sum, according to the argument presented in this section, there are particular moral facts, such as the fact *that your senior colleague's behavior is wrong*. These particular facts exist, however, in virtue of there being moral standards, such as the standard that it is wrong to humiliate others simply for pleasure. A close look reveals that standards such as this do not bear the marks of the natural. If things are as they appear, then that is good reason to believe that these standards are not natural.

The argument I've rehearsed in this section is subject to several objections. Let me close by voicing three reservations about it.

Earlier I noted that realists such as FitzPatrick hold that:

Moral standards are general moral facts.

While most realists would accept this claim, there are some who would deny it. For example, David Copp holds that moral standards are not facts but imperatives, which are neither true nor false; nonetheless, these imperatives ground particular moral facts. If Copp is right, the moral standard concerning humiliation mentioned earlier is really an imperative whose content is explicitly expressed by the sentence, "Don't humiliate others simply for pleasure!" (Copp 2007, 13–14).

While I do not wish to deny that imperatives can be used to express moral standards, I see no reason to believe that such standards are imperatives if realism is correct. Familiar considerations adduced against traditional forms of expressivism (which hold that the contents of moral judgments express pro or con attitudes and are not true or false in any robust sense) tell against such a view. It seems to make perfect sense, after all, to say such things as, "I know that it is wrong to humiliate others simply for fun," and, "It is a fact that it's wrong to humiliate others simply for fun," and so forth. Moreover, we seem to be able to embed propositions that have standards as their content in logical schemata such as modus ponens. If Copp's view were correct, however, then these sayings and embeddings would not make perfect sense. To make sense of them would require considerable work – work that realists have thought both unnecessary and ultimately futile (see Schroeder 2008). While not decisive, these considerations, I believe, weigh heavily against the view that standards are not facts but mere imperatives.

The second reservation I shall consider raises a different concern. When developing the argument in this section, I appealed to the claim that:

When a particular has a moral property, it has that property in virtue of there being some moral standard which that particular satisfies or fails to satisfy.

Some realists, however, reject this claim because they hold that there are no moral standards. Somewhat more specifically, these philosophers advocate a view known as *particularism*, which assumes that moral standards would have to be truths, and that there are no true and useful moral standards.⁷ At best, according to particularists, there are moral

7 Those familiar with the literature on particularism know that its advocates, such as Dancy (2004), present it in different ways. The description I offer of the view, I believe, captures its essence.

“rules of thumb.” These rules of thumb are as some think of Newton’s laws of physics: strictly speaking false, but helpful for getting around the world.

Why would particularists hold that there are no true and useful moral standards? The primary argument that these philosophers voice is one that appeals to the holism of reasons. Roughly stated, the idea is that a descriptive feature such as *being such as to humiliate simply for pleasure* is not *itself* a wrong-making property. Whether it is depends on contextual features. In some contexts, this descriptive feature provides powerful moral reasons not to humiliate; in other contexts, such as when we’re playing a game that involves humiliating other contestants, it does not. If this is so, we cannot claim that there are unqualified moral standards that are at once true and useful.

This last claim strikes me as correct. However, as best I can tell, absent other assumptions, little of interest follows from it. What nonparticularists should say is that most moral standards exist and are implicitly qualified by *ceteris paribus* clauses; these standards are, in ordinary conditions, true for the most part. Admittedly, specifying those conditions in which, say, humiliating merely for pleasure is not a wrong-making feature might not be easy, since those conditions would be fairly remote from ordinary experience. Nonetheless, this is compatible with these standards being true. And it would not follow that, if moral standards were qualified by *ceteris paribus*, they would be unhelpful when engaging in moral deliberations. To the contrary, they might be indispensable.

I turn to a final concern. Toward the beginning of our discussion, we distinguished between concepts, on the one hand, and properties and facts, on the other. Some (and perhaps all) propositions, I suggested, are composed of concepts, while some facts are composed of properties. Naturalists, I observed, might claim that while moral propositions are composed of non-natural moral concepts, these concepts characterize only natural properties or facts. Throughout our discussion in this last section, however, I have (mostly) spoken as if moral standards are facts. But one might deny this, claiming that such standards are propositions composed of non-natural moral concepts, while the properties to which they refer are natural. On the further assumption that such truths would hold in virtue of the essences of their constituent non-natural concepts (and not simply because of some further fact that makes them true), this view would imply that while there are non-natural moral truths (since they are composed of and made true by non-natural moral concepts), there are no non-natural moral facts. Instead, all facts are natural.⁸ This view, then, would deny FitzPatrick’s claim that:

Moral standards are moral facts.

In my view, this last objection is instructive for several reasons. For one thing, if it were correct, it would imply that the divide between certain versions of naturalism and non-naturalism might not be as wide as many have believed. After all, if the position we’re considering is correct, there is a thesis that both naturalists and non-naturalists can accept: namely, that there are non-natural moral truths.

Perhaps more importantly, however, the objection we’re considering illustrates that the argument offered for moral non-naturalism in this section is incomplete. It is true that we have uncovered no natural facts that are plausible candidates for being moral facts. Still, non-naturalists have more work to do if they are going to satisfactorily defend the claim

⁸ This is the position that Cuneo and Shafer-Landau (2014) call *minimal moral non-naturalism*.

that moral properties are non-natural and that moral standards are non-natural facts (as opposed to true propositions constituted by non-natural moral concepts). My own view is that non-naturalists can accomplish this work in a variety of ways, including by arguing that only facts can play the role of being that which explains why particulars have moral properties, as (2) asserts. However that may be, a benefit of our discussion is that we can better see exactly where more progress needs to be made.

Acknowledgments

Thanks to Kelly Clark and Russ Shafer-Landau for their comments on this chapter.

References

- Copp, D. (2007). *Morality in a Natural World*. Oxford: Oxford University Press.
- Cuneo, T. (2006). Moral Facts and Configuring Causes. *Pacific Philosophical Quarterly* 87: 141–162.
- Cuneo, T. (2007). Recent Faces of Moral Nonnaturalism. *Philosophy Compass* 2: 1–32.
- Cuneo, T. (2014). *Speech and Morality*. Oxford: Oxford University Press.
- Cuneo, T. and Shafer-Landau, R. (2014). The Moral Fixed Points: New Directions for Moral Nonnaturalism. *Philosophical Studies* 171(3): 399–443.
- Dancy, J. (2004). *Ethics Without Principles*. Oxford: Oxford University Press.
- Enoch, D. (2011). *Taking Morality Seriously*. Oxford: Oxford University Press.
- FitzPatrick, W. (2008). Robust Ethical Realism, Non-Naturalism, and Normativity. In *Oxford Studies in Metaethics*, vol. 3, edited by Russ Shafer-Landau, pp. 159–206. Oxford: Oxford University Press.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Gibbard, A. (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Gibbard, A. (2012). *Meaning and Normativity*. Oxford: Oxford University Press.
- Jackson, F. (1998). *From Metaphysics to Ethics*. Oxford: Oxford University Press.
- Jackson, F. (2012). On Ethical Naturalism and the Philosophy of Language. In *Ethical Naturalism: Current Debates*, edited by Susana Nuccetelli and Gary Seay, pp. 71–88. Cambridge: Cambridge University Press.
- Korsgaard, C. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- McPherson, T. (2012). Mark Schroeder’s Hypotheticalism: Agent Neutrality, Moral Epistemology, and Methodology. *Philosophical Studies* 157: 445–453.
- McPherson, T. (2015). What is at Stake in Debates among Normative Realists? *Noûs* 49(1): 123–146.
- Olson, J. (Forthcoming). The Metaphysics of Reasons. In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star. Oxford: Oxford University Press.
- Parfit, D. (2011). *On What Matters*, vol. II. Oxford: Oxford University Press.
- Plantinga, A. (2010). Naturalism, Theism, Obligation and Supervenience. *Faith and Philosophy* 27: 247–272.
- Railton, P. (1986). Moral Realism. *Philosophical Review* XCV: 163–207.
- Scanlon, T.M. (2014). *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, M. (2008). *Being For*. Oxford: Oxford University Press.
- Shafer-Landau, R. (2003). *Moral Realism*. Oxford: Oxford University Press.
- Sturgeon, N. (1988). Moral Explanations. In *Essays on Moral Realism*, edited by Geoffrey Sayre-McCord, pp. 229–255. Ithaca: Cornell University Press.
- Wedgwood, R. (2007). *The Nature of Normativity*. Oxford: Oxford University Press.