

# How to be a Moral Platonist

---

Draft for the 2013 Wisconsin Metaethics Workshop

Knut Olav Skarsaune

Supervenience relations are fertile ground for philosophical argument. There is often agreement that some properties A supervene on properties B, but disagreement about *why* this is the case. Friends of reduction say it is because the A-properties just are B-properties. Dualists say the two are distinct, but offer some explanation for why we nevertheless always find the A's in the company of their regular B-partners. A law of nature, perhaps, or interaction in the pineal gland. Still others try to find middle ground between reduction and dualism, so they can be “non-reductive physicalists” about the mind, or biology or chemistry.

We find a similar dialectic in metaethics, but with a further twist: most participants think it is not only true, but also *analytic*, that normative properties supervene on descriptive properties. Unlike the mental-physical or the biological-physical case, normative-descriptive supervenience seems to be enforced by the normative concepts themselves (more on this below). So metaethical theories have two things to explain: i) why normative properties supervene on descriptive properties, and ii) why this pattern is analytic.

These explanatory requirements are generally thought to have a special edge against the *nonnatural realist* view defended by G. E. Moore, Derek Parfit, Thomas Nagel and others. In the wake of two papers by Simon Blackburn (1971, 1985), an alleged inability of nonnatural realism to explain normative-descriptive supervenience has become one of the standard objections to the view. James Dreier, for example, thinks Blackburn has achieved “victory over Moorean dualism”:

Just what the connection [between value and natural properties] is, and in what sense moral properties “follow from” natural ones, Moore was never able to explain. The dualist, then, seems to be saddled with what Blackburn calls “an opaque, isolated, logical fact, for which no explanation can be proffered,” an extra law of metaphysics. (1992, p. 18)

Allan Gibbard again and again uses the issue to motivate expressivism over non-natural realism:

A non-naturalistic “moral realist” can present certain features of ethical concepts as brute truths: that, for example, whether an act is right or wrong depends on its natural properties. [...] Such a theorist, though, offers no explanation at all of the features of moral and other normative concepts. My aim in this book is to render normative concepts unmysterious, to explain those

features of ethical concepts that such a non-naturalist can only treat as brute. (2003, p 20)

My own theory explains much that non-naturalism takes as brute features of the non-natural realm. If the good exercises its own sovereignty, why does goodness depend on natural fact? That's just the way the concept works, the non-naturalist must be reduced to saying: it just does. (2003, p 184)

Michael Smith:

How can it be that *mere reflection* enables us to come by knowledge of which natural properties and which non-natural properties are coinstantiated, coinstantiated in a way that reflects the *a priori* supervenience of the moral on the natural? Given that, according to the non-naturalist, all we can say about non-natural properties *a priori* is that they are *simple* properties, neither constituted by nor analysable in terms of natural properties, it appears that they can give no answer. For them this must remain a mystery. (1994, p. 24)

*The Stanford Encyclopedia of Philosophy*:

[...] the supervenience of the moral on the non-moral seems to fall out of common sense rather than out of our scientific world-view. Since the methodology associated with non-naturalism itself puts great weight on common sense, the argument from supervenience seems like a promising one. At the very least, the argument from supervenience seems uniquely well situated to undermine non-naturalism without begging any central methodological questions. (2008)

These are in fact the last words of the entry on *Moral Non-Naturalism*.<sup>1</sup>

There is a good answer to this objection. In fact, we only have to make one small adjustment to G. E. Moore's account of normative concepts in order to explain why normative-descriptive supervenience is analytic. This idea was given to me some years ago by Kit Fine.<sup>2</sup> In a nutshell, the proposal is this: accept Moore's account of normative predicates as primitive *when they are applied to kinds*. But when they are applied to *particular* things, normative predicates express slightly different, shifted senses, which can be defined in terms of the kind-applying senses.

For example, let "good<sub>1</sub>" and "good<sub>2</sub>" express the particular-applying and the

---

<sup>1</sup> Written by Michael Ridge. See also Alexander Miller (2003, pp. 31-33), Graham Oddie (2005, p. 19), Dalia Drai (2000, pp. 27-28), Frank Jackson (1998, ch. 5).

<sup>2</sup> Written in Fine's barely legible hand on the back of a term paper. Many thanks to him for the idea, and for very stimulating conversations on the subject over the last few years.

kind-applying senses, respectively, of “good”. The former is definable in terms of the latter, as follows: by definition, a particular thing is good<sub>1</sub> if and only if it is a token of a good<sub>2</sub> kind.

Now if *a* and *b* are descriptively exactly similar, then they are tokens of exactly the same kinds, so *a* is a token of a good<sub>2</sub> kind if and only if *b* is. Therefore, by the definition, *a* is good<sub>1</sub> if and only if *b* is. Supervenience just falls out of the definition of “good<sub>1</sub>”.

This neo-Moorean account is a close cognitivist (truth-conditional) analogue of the analyses R. M. Hare (1952) and Allan Gibbard (1990) have offered over in the non-cognitivist camp. They also analyse ascriptions of a normative predicate to a particular in terms of a *general* commitment (in the case of Hare, a *universal prescription*, and in the case of Gibbard, endorsement of a *norm*). The neo-Moorean account implements this basic structure in a truth-conditional framework. In honour of Hare, I call it *cognitive universalism*.

I do not try to motivate cognitivism over non-cognitivism in this paper. My main message is that cognitivists of all stripes – even Moorean non-naturalists – can explain the analyticity of supervenience in the same elegant way that Hare and Gibbard have shown us.

I begin by looking at relevant passages in Moore and Blackburn (section 1). Section 2 gives a detailed presentation of cognitive universalism. My defence of this account has four legs. First, it enjoys direct intuitive support (section 2). Second, it nicely explains why normative-descriptive supervenience is analytic (section 2). Third, there is independent linguistic evidence for the existence of both particular-applying and kind-applying senses of the normative predicates (section 3). Fourth, it explains why moral epistemology has the structure it has (section 4).

In section 5, I briefly discuss the wider philosophical implications of the view. I suggest that cognitive universalism, as a semantic account, is compatible with anti-realist cognitivist views (such as *culture relativism* and *constructivism*), if these are construed as views about the nature of normative properties, rather than as analyses of normative language.

But I mainly focus on the implications for Moore’s descendants, the nonnatural realists. I argue that they should just repeat, at the metaphysical level, the explanation of supervenience that cognitive universalism provides at the conceptual level. Just as they should say that the basic normative concepts apply to kinds, so they should say that the basic normative *properties* take kinds, or perhaps properties, as their bearers. I argue, in other words, that nonnatural realists should be platonists.

## 1 Historical background

The natural place to start is with Moore's *Principia Ethica* (1903, henceforth *Principia*), because it is, in a sense, the null account: Moore argues that the basic normative concepts have *no* informative analysis. The concept he takes to be basic is *value*, or more accurately the concepts of positive value (*good*) and negative value (*bad*).<sup>3</sup> His main claim is that these are “simple and indefinable”:

‘good’ has no definition because it is simple and has no parts. It is one of those innumerable objects of thought which are themselves incapable of definition, because they are the ultimate terms by reference to which whatever is capable of definition must be defined. (*Principia*, pp. 9-10)

Since *good* cannot be defined, *a fortiori*, it cannot be defined in descriptive or naturalistic terms. This is why descriptions never entail evaluations. Take any description D that you like, it is always an “open question” whether some thing *x*, which is D, is also good. Likewise, the question “Is D-ness good?” is always open: we are never forced, on pain of conceptual confusion, to answer one way or the other (*Principia*, pp. 15-16).

Moore thinks the basic question in ethics is what *kinds* of things are good and bad (*Principia*, pp. iv, 118-120). I will argue later that this observation is absolutely crucial. But he does not draw a line from this point to the analysis of the concept *good*. He writes throughout as if *good* is “simple and indefinable” *also when applied to particulars*, that is, to particular people, acts and states of affairs, like Florence Nightingale or the Marshall plan. So, for example, the judgment that Ms. Nightingale was good applies a primitive, unanalysable concept *good* to that individual.

He also notes in passing that “a judgment which asserts that a thing is good in itself, [...] if true of one instance of the thing in question, is necessarily true of all; [...] all judgments of intrinsic value are in this sense universal [...]” (*Principia*, p. 27). But again he fails to draw any implications from this point to the analysis of the concept *good*.

He did, however, return to the issue later. In the 1922 paper *The Conception of Intrinsic Value* (henceforth *Conception*) Moore again grapples with the universalisability of normative judgments. He still holds to the view that the concept *good* is simple and indefinable, but he also sees that there limits to our freedom in how we apply the concept. Things have their value *in virtue of* what they are like descriptively, and therefore, there can be no evaluative difference where there is no descriptive difference:

---

<sup>3</sup> I use italics to refer to the concept expressed by the italicised word. I also use italics to refer to kinds. And I use them for emphasis. I try to write e.g. “the concept *good*” or “the kind *murder*” when context does not make it obvious which is intended.

When I say, with regard to any particular kind of value, that the question whether and in what degree anything possesses it *depends solely on the [descriptive] nature of the thing in question*, I mean [...] that if a given thing possesses [such a value] in a certain degree, then [...] anything *exactly like it* [descriptively must possess the value] in exactly the same degree. Or to put the point in the corresponding negative form: It is *impossible* that of two exactly similar things one should possess it and the other not. (*Conception*, 260-261)

This is as far as I know the first precise statement of normative-descriptive supervenience in the literature. But *Conception* ends in frustration: Moore agonises over what kind of necessity is involved in the statement he has just made. In what sense of “impossible” is it impossible for two exactly similar things to have different value? He admits he does not have a good answer.

It is not easy to interpret Moore’s vocabulary toward the end of *Conception*, but I believe the source of his troubles is this. He must have felt that the necessity involved was in some sense conceptual or analytic; that the concept of value does not allow for two descriptive twins to have different value. This is at least one way to understand his remark that the necessity is “unconditional” (*Conception*, p. 275). But at the same time, his commitment that *good* is simple and indefinable blocks him from saying what he wants to say. The necessity is “apparently” *not* “identical with the logical ‘must’” (*Ibid.*). For if *good* is simple and indefinable, why should it be incoherent to apply it to one but withhold it from the other of two descriptive twins? Where does this conceptual restriction come from?

Studying *Conception*, Simon Blackburn came to think that this issue generates a decisive objection to nonnatural realism:

Imagine a thing *A*, which has a certain set of naturalistic properties and relations. *A* also has a certain degree of moral worth: say, it is very good. This, according to the realist, reports the existence of a state of affairs, *A*’s goodness. Now the existence of this state of affairs is not entailed by *A* being as it is in all naturalistic respects. [...] That is, it is logically possible that *A* should be as it is in all naturalistic respects, yet this further state of affairs not exist.

[Now if some other thing *B* is just like *A* in all naturalistic respects], then it follows that *B* is also good. And this is a puzzle for the realist, because there is no reason at all, on his theory, why this should follow. If the goodness is, as it were, an *ex gratia* payment to *A*, one to which *A* is not as a matter of logic entitled in virtue of being as it is in all naturalistic respects, then it should be consistent to suppose that although goodness was given to *A*, it was not given to *B*, which merely shares the naturalistic features that do not entail the goodness. [...] Supervenience becomes, for the realist, an opaque, isolated logical fact for which no explanation can be proffered. (1973, in 1993, 118-9)

As we see, Blackburn takes as a premise that normative-descriptive supervenience is an analytic truth, in the sense that, for example

(1)  $a$  is good.

and

(2)  $b$  is descriptively exactly similar to  $a$ .

together conceptually entail

(3)  $b$  is good.

Like most commentators, I think Blackburn is right about that. But many philosophers are so suspicious of analytical truth that it may be worthwhile to dwell on the point for a moment. Here is Blackburn's comment:

One thing, then, that must be established in defending this part of the argument is that if somebody claimed, say, that an action was absolutely identical in every respect with another, except that it was much worse; or that a feature of character like courage had changed in no way in its nature, relations, consequences, but yet was of much less value than formerly; it would be a [conceptual] and not merely a moral mistake that had been made. (ibid, 116)

I will not attempt a general defence of analyticity in this paper. I can only report that, if I met someone who completely flouted supervenience, by calling some things "good" and other things "not good" while at the same time finding no descriptive difference between them, then I would be left with nothing else to think than that their word "good" meant something other than mine. It would be like talking with someone who said both " $x$  knows that  $p$ " and " $p$  is false".

At any rate, I hope it will be granted that there are no *special* reasons, peculiar to the case, to doubt the analyticity of normative-descriptive supervenience. It is analytic if anything is. And to the hardboiled analyticity sceptic, I would say this: pretend, for the duration of this paper, that you do believe in analyticity, and see if you do not like what we can *do* with it; see if you do not find the analysis of normative concepts that follows illuminating.

Let us proceed, then, on the assumption that (1) and (2) together conceptually entail (3). As we have seen, Moore's view cannot explain this entailment, and that

is Blackburn's objection.<sup>4</sup> It could be that his rhetoric is a bit rich; after all, it is not unheard of for a philosophical theory to leave some aspect of its subject matter as basic and unexplained. As David Lewis remarked in another context, a theorist faced with a putative explanandum always has three options. She can say: "I deny it", or "I explain it thus...", or she can say "I accept it as primitive" (1983, 352).

If one takes the third option, the question is how philosophically embarrassing it would be to leave this particular explanandum unexplained. We have seen that James Dreier, Allan Gibbard, Michael Smith, the *Stanford Encyclopedia of Philosophy* and many others think it would be highly embarrassing. But we will not need to haggle over the price. For it is in any case a mistake to take something as primitive when it *has* an explanation. And that, I will argue, is the case here.

## 2 Of particulars and kinds.

Let us return to Moore's claim that the concept *good* is "simple and indefinable". Is that always true?

Suppose your moral guru tells you that Bob did something good yesterday. You trust your guru completely, so you form the belief that Bob did something good yesterday, even though you have no idea what he was up to. Now does it seem right that the concept *good*, as it figures in this belief, is simple and indefinable?

---

<sup>4</sup> Blackburn later (1985) restated the argument as follows: Suppose we judge a thing to be good, on the basis of its having certain descriptive properties and relations. Put all these together in a big, conjunctive property F. Include in F also all its *negative* properties, that is, include *not being G*, if being G would have destroyed its goodness. The big descriptive property F, then, *suffices* for goodness, or so we think. Now, since they think descriptions never entail evaluations, Mooreans will think there are *conceptually possible worlds* in which the F's are not good. In addition, of course, they will think there are conceptually possible worlds in which the F's are good. But since, by hypothesis, *every* normatively relevant property and relation, both positive and negative, are included in F, they will deny that there are any conceptually possible "mixed worlds", in which some of the F's are good and others not.

This strikes Blackburn as odd. If there are conceptually possible worlds in which the F's are good, and ones in which they are not good, why are there no conceptually possible worlds in which, say, half of them are good? Blackburn thinks it is implausible to have such a "ban on mixed worlds" without giving any explanation for it.

The account I will offer meets this restated argument in exactly the same way as it meets the original version, so I will not discuss it any further in the text.

It seems to me that your belief can in fact be analysed. What you believe is not that Bob performed some act that simply had the property of being good, period, end of story. You believe that he performed an act with *some descriptive properties or other* that made it good. And it follows from what you believe that anyone who acts *likewise* in *like* circumstances will also be doing something good. In short, what you believe is that Bob *acted in some way or other* such that *acting that way is good*.

This is the intuitive starting point of the account I will propose. As usual with intuitive starting points, I do not have much to say by way of argument to support it. It is based on introspecting this structure in my own beliefs. But search your mind, and you will see that it is so.

Nor do I want to say that the account I will propose is the *only* way to go from this intuitive starting point. The details of the account will be supported by other, linguistic arguments. But it will be convenient to put the account on the table first, and give the supporting arguments afterward.

The proposal is this: accept Moore's account of normative predicates as primitive *when they are applied to kinds*. But when they are applied to particulars, normative predicates express slightly different, shifted senses, which can be defined in terms of the kind-applying senses. So for example, the word "good" can express two different, but closely related senses, depending upon whether it is applied to a kind or to a particular. The particular-applying sense, *good*<sub>1</sub>, can be defined in terms of the kind-applying sense, *good*<sub>2</sub>, as follows:

CU:  $\text{good}_1(x) \leftrightarrow_{\text{def}} \exists K [\text{token}(x, K) \ \& \ \text{good}_2(K)]$

In other words,  $x$  is *good*<sub>1</sub> iff there is a kind  $K$  such that  $x$  is a token of  $K$ , and  $K$  is *good*<sub>2</sub>.

The variable " $K$ " ranges over descriptive kinds. For example, *war* is an event kind; *eating bananas* is an act kind, *being happy* is a kind of mental state. There is no restriction on how general or specific the kinds can be, so for example, *eating bananas while sitting on a train passing by a lake* is also in the range of  $K$ .

However, we should impose the following restriction. We do not include so-called *haecceitic* kinds in the range of  $K$ . A kind is haecceitic if it concerns a specified individual. So, for example, *buying Mary a bucket of roses* and *moving to Dallas* are haecceitic kinds. The motivation for this restriction is that normative concepts do not permit mere haecceitic differences to make a normative difference.

Some writers deny this point. Amongst other examples, Matthew Kramer (2009, ch. 10) points out that many religious believers treat being pleasing to God as normatively relevant, but they can hardly be accused of conceptual confusion. He also imagines a person who prefers acts that benefit France, but not because France has any interesting descriptive properties, or because he is French or stands in any other interesting relation to France.

I do not find these examples convincing. The example from religion is misleading, because actual religious believers think God's opinions matter because he has certain qualitative properties (e.g. power, wisdom) and stands in certain qualitative relations to them (has created them, cares about them).<sup>5</sup> It is far from clear that any actual religious believers think God's haecceitic identity makes a difference, so that, for instance, we could have another qualitatively identical world, with a qualitatively identical creator, but the creatures over there have no reason to obey *their* creator.

The thought experiment with the devoted Francophile is also weak. What we need to imagine is a situation where we have two qualitatively identical countries, France and France\*, say, and a person who stands in the same qualitative relations to both. So he did not, for example, spend the summers of his youth in one of them; nor does he have different feelings towards them. He knows all this, but still, on the basis of no other difference whatsoever, he calls act that benefit France "good" and acts that benefit France\* "not good". Then I would simply repeat the point from section 1: faced with such a person, I would be left with nothing else to think than that his word "good" meant something other than mine.

The following is at stake. In section 1, I assumed that normative-descriptive supervenience is analytic. There is a pretty firm consensus in the literature about that. Now we are considering something stronger: whether it is analytic that normative properties supervene on *qualitative* descriptive properties. I say it is, and my sense is that that is what most participants in the debate have also meant. But my argument does not hinge on the point: if the reader thinks normative concepts do allow haecceitic properties to make a normative difference, she can suit the account to her liking by simply allowing haecceitic kinds in the range of the kind-variable *K* in the definition.

Like properties, kinds can have instances/tokens. So, for example, the Thirty-years' war is a token of the kind *war*, and I am an instance of the kind *Homo Sapiens*. A kind and a particular stand in the *tokening-relation* just in case the particular is an instance of the kind.

The right-to-left direction of CU predicts that, if *good*<sup>2</sup> applies to a kind, then *good*<sup>1</sup> applies to *every* instance of the kind. That might seem too strong. For example, we might be inclined to accept both

---

<sup>5</sup> "Qualitative property" is hard to define. For my purposes, we can understand it as "non-haecceitic property". It is important to many religious believers that God is the *unique* possessor of certain properties (creator of the world, say); perhaps also that God possesses certain properties that can at most be had by one subject, like being almighty. But these remain non-haecceitic properties, in the relevant sense, because they can be defined without the use of a name or a demonstrative.

- (4) Knowledge is good.

and

- (5) Oedipus' learning that Jocasta was his mother was not good.

But this is not a counterexample. The explanation is that “good” has two senses along another dimension. It is generally agreed that normative predicates can express both *pro tanto* and an *all-things-considered* senses.<sup>6</sup> And we must distinguish these *before* we apply CU. So, for example, in (4) we have *pro-tanto-good*<sup>2</sup>, and in (5) we have *all-things-considered good*<sup>1</sup>. Under those readings, (4) and (5) are compatible with CU.

This is not an *ad hoc* move: the distinction between the *pro-tanto* and *all-things-considered* senses of the normative predicates is independently motivated. And the prediction that, if a normative concept applies to a kind, then the corresponding particular-applying concept applies to every instance, is confirmed by intuition. For example, if you think (5) is true, then you will not get (4) to be true if you force yourself to read it in the *all-things-considered* sense. And if you think knowledge is *pro tanto* good, you will not get (5) to be true if you force yourself to read it under the *pro tanto* sense.

So far I have discussed “good”, but I think parallel analyses apply at least for “bad”, “right”, “wrong”, “just” and “unjust”, or more accurately for both the *pro tanto* and the *all-things-considered* senses of these.<sup>7</sup> CU, then, works as a general recipe for defining particular-applying senses of normative predicates in terms of their kind-applying senses.

This account is both original and not. The core idea is looted from R. M. Hare's *universal prescriptivism* (1952), and Alan Gibbard's *norm-expressivism* (1990). These accounts also analyse ascriptions of a normative predicate to a particular in terms of a *general* commitment; in the case of Hare, a universal prescription, and in the case of Gibbard, endorsement of a norm. In both cases, the general com-

---

<sup>6</sup> If an act is *pro tanto* wrong, then it has something wrong about it, even though it may also have something right about it. If it has more wrong than right about it, then it is *all-things-considered-wrong*. Some writers distinguish *pro tanto* concepts from *prima facie* concepts (the difference is not important for my purposes here); if both exist, then we can apply CU to both versions.

<sup>7</sup> I think analyses similar in spirit, but perhaps different in detail, apply for “reason” and “ought” as well. I ignore the question of whether some of these normative concepts can be defined in terms of the others.

mitment is existentially quantified.<sup>8</sup> What the new account achieves is to make that insight of Hare and Gibbard's available in a truth-conditional framework. In a nod to Hare, we can call it *cognitive universalism*.

I like to think the philosophical allure of cognitive universalism is immediate and irresistible. But just in case, I will now proceed to give some arguments in its favour. The first "argument" is just the motivating intuition described above: that normative claims about particulars always seem to have a *general* content; they say that the particular has *some* set of descriptive properties in virtue of which it, and anything else that might share the same descriptive properties, has the value it has.

Secondly, cognitive universalism can easily explain why normative-descriptive supervenience is an analytic truth, and thereby answer Blackburn's objection. For example, we had

- (1) *a* is good
- (2) *b* is descriptively exactly similar to *a*.

The explanation goes as follows. From (1), by UC, it follows that there is some kind, let us call it "L", such that *a* is a token of L, and L is good<sub>2</sub>. From (2) and our definitions of kind and of the tokening-relation, it follows that *b* is a token of exactly the same kinds as *a*. So in particular, *b* is a token of L, which, recall, is good<sub>2</sub>, and so by UC,

- (3) *b* is good.

It should be straightforward to see how this generalises to the other normative predicates, and also to the case where *a* is not good (in which case it follows that *b* is not good either).

Relative to the Moore of *Conception*, we have replaced a situation where we had a primitive concept (*good*), and a brute conceptual necessity (supervenience), with a situation where we have a primitive concept (*good*<sub>2</sub>) and a defined concept (*good*<sub>1</sub>). Supervenience just falls out of the definition of *good*<sub>1</sub>.

These are the two main attractions of cognitive universalism. Now the view also has a couple of features that may be unfamiliar or off-putting to some readers. First, it treats general normative claims, like "Knowledge is good", as kind-referring, and posits suitable kind-applying senses of the normative predicates. In my experience, many philosophers are uncomfortable with this: is it not more plausible to treat general normative claims as quantificational? Is there any lin-

---

<sup>8</sup> For Gibbard, to say that a particular act is rational is to express acceptance of *some* norm that permits it. For Hare, to say that a particular act is right is to command everyone to act *likewise* in *like* circumstances (without specifying the act or the circumstances).

guistic evidence (beyond supervenience) to suggest that these kind-applying senses even exist?

In the next section, I defend the answers No and Yes, respectively, to these questions. My strategy will be to apply, to normative language, a line of argument in favour of kind-reference familiar from the literature on *generics* like “Bears hibernate” and “Dogs have four legs” (see Carlson 1977, Liebesman 2011).

However, and secondly, it is important to see that cognitive universalism is not just a routine application of the kind-reference approach to generics. It takes a radical further step by reversing the usual order of conceptual priority between the kind-applying and particular-applying senses of predicates. Usually, on the kind-reference view, we will think the particular-applying senses are primary, and the kind-applying senses are derived from them. So for example, when we do the lexical semantics of “have four legs”, as applied to dog-kind, we find that it expresses the concept of being such that one’s normally developed, unmolested tokens *have four legs in the particular-applying sense*. So in the usual case, *after we have done lexical semantics*, we will end up with truth conditions concerning particulars and their properties. But cognitive universalism finds the opposite pattern in normative language; it takes the kind-applying senses as basic and the particular-applying senses as derived. This means that, *even after we have done lexical semantics*, we will be left with truth conditions concerning kinds. This is a much more radical view than the standard kind-reference view about generics. Backing it up will be the job of section 4.

### 3 Why we need kinds in our semantics anyway

So far, I have explained cognitive universalism as the combination of two views. The first is the CU-analysis, defining particular-applying normative concepts (*wrong<sub>1</sub>*, etc.) in terms of kind-applying concepts (*wrong<sub>2</sub>*, etc.). The second is that these kind-applying concepts are in turn directly expressed in general normative claims, like “Tax evasion is wrong”.

These two parts of the view are logically independent. In theory, one could accept the CU-analysis, but maintain that normative predicates invariably express what I call their “particular-applying senses”, both in particular and general claims. On this view, general claims like “Tax evasion is wrong” would instead be submitted to some form of *quantificational* analysis, such as “For all  $x$ , if  $x$  is an instance of tax evasion, then  $x$  is wrong”. That way, since the predicate “wrong” is applied to a variable that ranges over particulars, it would still express the same, particular-applying sense. As for *wrong<sub>2</sub>* and the other kind-applying concepts posited in CU, they would as it were be “silent”; they would figure as *parts* of the concepts we express, but would never be articulated on their own in natural language. Still, the explanation of supervenience would go through as before.

So why not take this alternative view, which offers the same explanatory payoff, but without committing to the kind-reference view of general normative claims? Primarily, I will argue, because the kind-reference view is independently plausible. But also because it lends support to the CU-analysis in a couple of ways. The first is simply that it is more plausible to posit a concept, and claim that it plays an important role in our thought, if one can actually point to examples where it is expressed in natural language. If questions of which particulars are good<sub>1</sub> reduce to questions of which kinds are good<sub>2</sub>, then you would expect that we sometimes think about which kinds are good<sub>2</sub>, and that we occasionally feel the urge to communicate such thoughts to other people. At the very least, a defender of the alternative view just mentioned will have some explaining to do.

The second way the kind-reference view supports the CU-analysis concerns acquisition. Simply put, if *wrong*<sub>1</sub> is analysed in terms of *wrong*<sub>2</sub>, and *wrong*<sub>2</sub> is never articulated, then how do children learn *wrong*<sub>1</sub>? If one concept is analysed in terms of another, we would expect learning the latter to be a part of the story about how one learns the former. For example, in the normal run of things, learning the concept *brother* is a part of the story about how we learn *uncle*. I do not intend any strong impossibility claim here: Perhaps it is somehow possible to learn *uncle* without learning *brother*. Or perhaps there is some relevant disanalogy between the *uncle/brother* case and the *wrong*<sub>1</sub>/*wrong*<sub>2</sub> case. But again, a defender of the alternative view will at least have some explaining to do.

On the other hand, we get a very attractive acquisition story when we combine the CU-analysis with the kind-reference view. For notice how natural it is, when praising or reprimanding a child, to emphasise the normatively relevant kind its behaviour exemplified. “Stealing cookies is bad”, we say. “Did you let your sister go first? Very good!”. Such exchanges introduce children to the idea that certain kinds of act, certain ways of acting, have normative significance. And the idea that a *particular* act is wrong, good, etc., if and because it is an act of one of these kinds seems to follow naturally.

If claims like “Stuffing noodles up your brother’s nose is wrong” are kind-referring, then, we have independent confirmation that *wrong*<sub>2</sub> exists, and we have a good explanation of how children learn it. But kind-reference analyses are quite controversial. Ever since Frege taught us to use quantifiers to replace Aristotle’s baroque logic of syllogisms with predicate logic, and Russell used them to retire the king of France, philosophers have learned to love quantifiers, and to use them to analyse general claims whenever possible. Thus, for example, Donald Davidson:

Confusion over the relation between ordinary sentences about actions, and particular actions, has led some philosophers to suppose that these sentences are about *generic* actions, or *kinds* of actions. [...] Analogous remarks apply to the idea that ‘Lying is wrong’ is about a kind of action. ‘Lying is wrong’ may be rendered, ‘For all *x* if *x* is a lie then *x* is wrong’ or even, ‘The class of lies is

included in the class of wrong actions', but neither of these says that a kind of actions is wrong, but rather that each action of a kind is wrong. (1980, 168)

In conversation, I have found that many philosophers share Davidson's attitude, at least so far as to consider quantificational treatment the default option, placing the burden of proof on the kind-reference view.

I wish I had a knock-down argument to refute any quantificational treatment of claims like "Stealing cookies is wrong". Unfortunately, I do not. But there is an argument available, which seems like a knock-down to me, for a weaker conclusion: namely that *a certain kind* of general normative claim is kind-referring. And once that is established, considerations of uniformity push in the direction of a kind-reference account also of central cases of the form "To  $\phi$  is wrong/good/etc." and " $\phi$ -ing is wrong/good/etc.".

Before I give this argument, let us note that Davidson's proposal, as it stands, is a non-starter. For if we apply it to sentences about act kinds that lack actual instances, the result will be vacuous truth no matter what the predicate is. For example, suppose that no one is ever tortured for fifty years. Then, the claim "Torturing someone for fifty years is good" will come out true on Davidson's account, since

- (6) For all  $x$ , if  $x$  is a fifty-year-long-torturing, then  $x$  is good.

is vacuously true. But this problem is easy to fix. In my experience, people sympathetic with the account invariably respond by offering

- (7) Necessarily, for all  $x$ , if  $x$  is a fifty-year-long-torturing, then  $x$  is good.

as their analysis instead. Let us call this amended view *the Davidsonian account*.

I think this amended proposal is also wrong. But before I criticise it, it will be convenient to look at another quantificational approach. I will then criticise these two accounts together. The other account assimilates normative claims about kinds to so-called generic claims, or *generics*, like

- (8) Dogs bark.  
 (9) *Canis Familiaris* has four legs.  
 (10) The dog's gestation lasts 65 days.

Furthermore, the account analyses generics in terms of a *generic quantifier*, GEN, like so:

- (8\*) GEN $x$  (Dog( $x$ )) [barks( $x$ )]

This says roughly that, in general, if  $x$  is a dog then  $x$  barks. The precise semantic behaviour of GEN is a matter of disagreement. What is clear, however, is that it cannot simply dictate that some specified proportion, for example more than half, of the kind at issue (dogs) should satisfy the predicate (bark). This is because

the required proportion seems to vary from case to case. So, for instance,

- (11) Birds lay eggs.

is true even though less than half of all birds ever lay eggs, and

- (12) Sharks attack bathers.

is true even though only a tiny minority of sharks ever do this. But “Dogs are fine swimmers” would be false if a similar proportion of dogs were fine swimmers. In general, the proportion required seems to be lower the more shocking and striking the predicate. The upshot is that GEN must somehow “look inside” the sentence, at the predicate, in order to “know” which proportion to require.

There is an immediate contrast between normative claims about kinds, and *characterising generics*, like (8) - (10). A diagnostic test for characterising generics is that inserting “Usually/normally/generally” will result in at most a minor shift in meaning (Krifka *et al* 1995, 9). For example, we do not get significant shifts in meaning here:

- (13) Dogs have four legs.

- (13') Normally, dogs have four legs.

- (14) Running is fun.

- (14') Usually, running is fun.

But if we try this test on normative claims, we get a clear shift in meaning.

- (15) Firing a secretary because he refuses to sleep with you is wrong.

- (15') Usually, firing a secretary because he refuses to sleep with you is wrong.

- (16) Abortion is wrong.

- (16') Normally, abortion is wrong.

But this contrast does not in itself refute the GEN account of normative claims. Remember that we found that GEN must “look inside” the sentence in order to know what proportion to require. What a defender of the GEN account should take from these examples is that, when GEN “sees” a normative predicate, it imposes a very strict proportion requirement, perhaps approaching 100%. That is why, she should say, we get a weakening when we insert “usually/normally/generally”.

In fact, she should also consider uninstantiated kinds (say, fifty year long torturings), and conclude that, when GEN sees a normative predicate, it even imposes a *modal* distribution requirement; perhaps approaching 100% of the cases in 100% of the worlds. So we see that the most plausible development of the GEN account, in application to normative claims about kinds, will be substantively

very similar to the Davidsonian account (which, recall, used necessitated universal quantification). I will therefore criticise the two accounts together.

The argument I will give is an application, to normative language, of a strategy used by Greg Carlson (1977) and David Liebesman (2011) to defend the kind-reference approach to generics. The first step of my argument is the observation that act-kind noun phrases can occur together with so-called *kind-selecting predicates*, such as *invent*, *widespread*, *extinct*. Witness:

(17) Skateboarding was invented by bored surfers.

(18) Cain invented fratricide.

Why are these predicates called “kind-selecting”? The reason is that they are true of kinds, but are not true of any of their instances. For example, “invented by Cain” does not apply to any of the particular fratricides.

Granted, we can come up with a more complicated quantificational story that captures the truth conditions (roughly, Cain committed a fratricide at some time *t*, and there were no fratricides before *t*). Similarly, maybe it would be possible, with sufficient ingenuity, to come up with quantificational truth conditions for

(19) Mehmed legalised fratricide.<sup>9</sup>

(I will leave this as an exercise for the reader; what makes it tricky is that (19) could be true even if there never were any fratricides, but we do not want it to be *vacuously* true in that case.)

If we multiply examples like these, a defender of the quantificational approach would have to conjure up ever more complicated hypotheses about logical form. That is not in itself a damning objection – in one sense, these claims do have complicated meanings. But the problem is that the complexity does not seem to lie with the *sentences*, but rather with the *predicates*. We will capture what is going on better if we assign sentences like (19) a rather simple logical form, and instead assign “legalise” a rich lexical semantics.

This point is reinforced by the fact that the inference from (18) and (19) to

(20) Mehmed legalised something Cain invented.

seems like a very *simple* inference. But on the quantificational approach it would not be, because all these sentences would have hugely complex logical forms.

The natural view to take of these examples is that, *as far as sentence semantics is concerned*, noun phrases occurring together with kind-selecting predicates really do refer to act kinds. So, for example, the logical form of (19) and (20) are just

(19\*) legalised(*Mehmed*, *fratricide*)

---

<sup>9</sup> Mehmed II legalised fratricide in the Ottoman royal family in certain situations, in order to prevent wars of succession.

(20\*)  $\exists x$  [legalised(*Mehmed*,  $x$ ) & invented(*Cain*,  $x$ )]

This view does not have the silly consequence that (19) is true because Mehmed performed some mysterious act on an abstract object (an act kind, fratricide). Clearly, what *ultimately* makes (19) true is that he performed certain particular acts; perhaps he signed a document or made an oral announcement in the town square. Similarly, (17) is true because a gang of bored surfers performed some particular, pioneering acts of skateboarding. That is how you legalise, or invent, an act kind. But the point is that this information belongs in the lexical semantics of the predicates, not in the logical form of the sentences. In other words; understanding what these claims ultimately demand about the occurrence of particular acts in the world is not a matter of understanding the logical form of the sentences, but of understanding what the predicates mean.

The upshot is that we must distinguish between two levels of truth conditions. First we have the results of semantic analysis of sentences; we can think of the results of this analysis as *sentence truth conditions* or *logical form*. But it would be a mistake to think of these as giving us a picture of what the world must be like in order for the sentence to be true. We must check with lexical semantics first, to see whether any of the items that occur in logical form should be further analysed, before we get to what we can call *worldly truth conditions*.<sup>10</sup> So for example, because we understand what “widespread” means, we understand that the worldly truth condition of “Bedbugs are widespread in Shanghai” does not involve some abstract object being spread out over Shanghai, but rather that there be concrete bedbugs at many locations across the city.

With this two-step methodology in hand, then, we should conclude that, at the level of logical form, act-kind noun phrases really do refer to act kinds, when they occur together with kind-selecting predicates. But then we are forced to accept that the same thing sometimes happens under normative predicates, namely when normative and kind-selecting predicates are conjoined:

---

<sup>10</sup> Related distinctions are drawn, in a similar context, by Koslicki (1999) and Leslie (2008).

The division of labour I am describing here, between sentence semantics and lexical semantics, should not be confused with the division of labour between semantics (as a whole) and metaphysics. How to draw *that* line is another huge subject. But I take it that there are clear cases on either side of the line. For example, understanding what kind-selecting predicates demand about particular tokens is clearly on the semantic side. Someone who does not understand that “The Dodo is extinct” is true iff there once were Dodos, but none now, simply does not understand what “extinct” means. On the other hand, reduction of truths about chairs to truths about atoms arranged chairwise is clearly on the metaphysical side. One cannot get this reduction just on the basis of understanding concepts.

(21) Murder is widespread but wrong.

The conclusion of this argument, then, is that “murder” sometimes refers to a kind, and that “wrong” sometimes expresses a sense that applies to kinds. This conclusion generalises to other act-kind noun phrases and normative predicates (just run the argument again with other examples).

Even though cases like (21) are pretty rare, the point remains that they are nearly undeniable examples of kind-applying normative concepts being expressed in English. This alone gives a measure of support to the CU analysis – it is independent confirmation that we understand concepts that are suitable to figure on the right hand side of CU.

But like Carlson, Liebesman and others, I also think of these sentences as the thin end of a wedge that can be used to defend kind-reference analyses of a wider range of cases. Simply put, the further argument is that we have should give a uniform treatment of, say

(22) Tax evasion is wrong.

(23) Tax evasion is widespread.

(24) Tax evasion is widespread and wrong.

For if (22) is quantificational but (23) and (24) are kind-referring, then why can we infer (24) so easily from (22) and (23)? Unless we treat “Tax evasion” as kind-referring in all three sentences, this inference will either be invalid or, at best, enthymematic (with a suppressed intermediate inference from  $GENx (Tax\ evasion(x)) [Wrong(x)]$  to  $Wrong(tax\ evasion)$ ). But it sure *seems* like a simple and complete inference.

Together, these two arguments are the crux of the positive case in favour of the kind-reference approach to central cases like “To  $\phi$  is wrong/good/etc.” and “ $\phi$ -ing is wrong/good/etc.” (and indeed to generics in general). To have a full-blown defence of the view, one will also need to do a lot of defensive argumentation, providing alternative explanations for the kinds of data that have attracted people to the GEN approach. But on that front, I have little to add to Liebesman (2011), so I will simply refer the reader to that fine paper.

Let me sum up. I have suggested that, even though they are logically independent, the CU-analysis is most plausible when combined with a kind-reference analysis of general normative claims, because we can then say that the kind-applying concepts posited in CU are articulated in natural language, and we have an explanation of how children learn them. I then proceeded to give an argument in favour of the kind-reference view familiar from the literature on generics. A crucial move in (my version of) this argument was to distinguish between two levels of truth conditions: *logical form*, which concerns the interpretation of sentences, and *worldly truth conditions*, which also take lexical semantics into account.

With this distinction available I argued that, at the level of logical form, kind-selecting predicates are applied to kinds. But then considerations of uniformity favour a kind-reference account also of ordinary generics, and of general normative claims.

#### 4 Why normative judgments about kinds are basic

There remains, though, a deep contrast between normative language, on the one hand, and at least most descriptive language, on the other. Even if kind-reference is fairly common at the level of sentence semantics or logical form, in the descriptive case it will at least usually be “analysed away” in lexical semantics, so that, when we get to worldly truth conditions, we are left with only particulars and their properties and relations.

For example, in the case of kind-selecting predicates like “invent” and “widespread”, even though these are applied to kinds in logical form, this kind-reference is analysed away in lexical semantics. Just by knowing what “widespread” means, we know that what it is for a kind to be widespread is for it to have instances in many, scattered locations. So the worldly truth conditions of “widespread”-claims are about particulars and their properties and relations.

There may also be areas of descriptive language that do not work like this; areas where kind-reference and predication of properties to kinds persist all the way to worldly truth conditions. I take no stand on that issue. But I will argue that normative language works in this second way. In fact, normative worldly truth conditions are *always* about kinds, and *never* about particulars.

Let me explain. Cognitive universalism takes the kind-applying senses of normative predicates as basic. So a normative claim about a kind, like “Causing needless pain is wrong”, will not be reduced in terms of particulars at any level of semantic analysis. Its worldly truth condition is just that acting a certain way, *causing needless pain*, is wrong.

For normative claims about particulars, like “The Rwanda massacre was wrong”, the situation is a bit more complicated. Cognitive universalism defines the particular-applying senses of normative predicates in terms of the kind-applying senses. So application of normative concepts to particulars is analysed away at the level of lexical semantics. However, they are given *mixed* worldly truth conditions, in part descriptive, in part normative. In this example, the worldly truth condition is that the Rwanda massacre was a token of some kind or other, such that that kind is wrong. So the *descriptive* part of the truth condition is still about the particular, the Rwanda massacre. But the *normative* part, saying that the relevant kind is wrong, ascribes a normative property to a kind, not to a particular.

Cognitive universalism, then, predicts that normative judgments about kinds are

*basic*, in the following sense: the worldly truth conditions of normative claims are either purely about kinds, or else mixed, but with the normative part being about kinds. Normative language *bottoms out* in truth conditions about kinds. I do not want to hide the fact that this is a radical view.<sup>11</sup> But in this section, I explain why I think it is true.

A lot of the motivation has been given already. I have reported a basic intuition, namely that, when we judge that a particular act or thing is wrong or good, what we judge is that the thing or act is some way or other, descriptively, such that acting that way or being that way is wrong or good. Cognitive universalism accommodates this intuition. And it explains why normative-descriptive supervenience is analytic. We could accept the view on these grounds, and just take onboard as a consequence that normative language bottoms out in truth conditions about kinds.

But we can also give an independent argument for this aspect of the view. It is plausible that normative claims about kinds are in this way basic, because they are *epistemologically* basic in a parallel way. Normative epistemology also bottoms out in judgments about kinds, as I will now explain.

Let us start with a simple observation. In descriptive enquiry, we typically go from judgments about particulars to judgments about kinds. So, for example, we might do ornithology in roughly the following fashion:

- (A) This bird sings in the morning and that bird sings in the morning and yonder bird... – and come to think of it, they are all robins! So it seems robins sing in the morning.

But notice how backwards it would be to try to do ethics in a similar way:

- (B) This act is wrong and that act is wrong and yonder act... – and *by golly*, they are all sexual harassments! So it seems sexual harassment is wrong.

That is just silly. The direction of epistemic justification is the reverse in the normative case: we go from *general* normative judgments, and *empirical* judgments about particulars, to normative judgments about those particulars. For example, we go from the general judgment that using a position of power to pressure someone into bed is wrong, and the empirical judgment that that's what *x* did to *y*, to the judgment that what *x* did to *y* was wrong.

Granted, there are complications. We do sometimes arrive at normative judgments about kinds through investigation of their instances. For example, we probably do not have a direct intuition that appointing relatives to government jobs is wrong. We arrive at this judgment by looking at countries where that practice is common, and observing its typical effects. But this is not a counterexample

---

<sup>11</sup> But notice, again, the similarity with Hare and Gibbard's views.

to the point I am making. For what we then investigate about these instances is not their normative, but their *descriptive* properties. For example, we see that they lead to inequality of opportunity and loss of general utility. But our judgment that general utility and equal opportunity are good are in turn judgments about kinds, and *not* based on normative judgments about particulars.

The extreme case of this epistemic structure is pure hedonic utilitarianism. For the committed utilitarian, normative enquiry will be a lot like empirical enquiry; particular cases will be decided by investigating their effects, and judgments about kinds will be a lot like inductive generalisations. But at the bottom of all this epistemic activity is the single ur-judgment that an act is right if and only if it maximises the balance of pleasure over pain. This is a judgment about kinds, and is *not* in turn arrived at by inference from particular cases. Take it away and the whole superstructure evaporates.

The claim, then, is that all normative justification *bottoms out* in judgments about kinds. That may seem to contradict the popular view that normative epistemology proceeds by the method of *reflective equilibrium*; going back and forth between principles and cases, trying to find intuitively attractive principles that yield intuitively attractive verdicts about cases, revising both kinds of intuitions as we go. But not really. For notice that “cases” here are not really *particulars*; they are instead narrowly circumscribed kinds. For example, we might well have as a fixed point in our normative reasoning that the Rwanda Massacre was wrong. But this fixed point is not really that a particular act *de re*, the Rwanda Massacre, was wrong. For imagine we were to find out that the event was in fact not a mass murder, but a desperate hospital program trying to help the sufferers of some strange contagious illness. That would make us revise our normative beliefs about the event, and perhaps about particular political leaders, etc. But these “local” changes would not ramify into our wider normative belief system. We would still believe that *killing six hundred thousand people on account of their ethnicity* is wrong, and this belief would continue to play the role in our reflective equilibrium that we had hitherto (misleadingly) ascribed to the belief that the Rwanda Massacre was wrong.

The same point applies to everyday uses of our moral sensibility. Suppose for example that you witness a man subjecting another to some kind of humiliating treatment for no good reason. You have an immediate gut reaction telling you that what the first man is doing is wrong, and this in turn leads you infer that it is wrong to humiliate others needlessly. One might think that here, surely, is a case where a normative verdict about a particular case supports a general normative conclusion. But that would again be misleading. For suppose you learn that the two men were in fact stage actors rehearsing a play. That would lead you to revise your normative verdict about that particular act (*de re*). But it would not lead you to conclude that it is OK to humiliate others needlessly after all; the lesson you

learned about *that* still holds good.<sup>12</sup>

In general, what matters to your normative thinking is not really your reaction to particular cases *de re*, but your reaction to the descriptive properties you *think* these cases have, that is, to the kind you think they instantiate. What particular cases can do is to make this or that kind *salient* to us, by making a token salient, but it is our verdict about the kind that plays a role in our reflective equilibrium, not our verdict about the token. In this way, beliefs about particulars are *epiphenomenal* in our normative belief system; they are supported by but do not support beliefs about kinds. Reflective equilibrium is reached when our beliefs about more general kinds fit with our beliefs about more specific kinds.

Normative enquiry, then, has exactly the structure we would expect if cognitive universalism is true. Since normative claims about kinds are not generalisations over their instances, they are not justified in the way generalisations are justified, from premises about instances. Instead, they are justified either by some kind of direct conviction about the kind (say, that it is wrong to kill people on account of their ethnicity), or else by their coherence in a network of judgments about kinds. And just as normative claims about particulars, in worldly truth conditions, factor into a *descriptive* component about the particular, and a *general* normative component, so they are justified by evidence about the descriptive properties of the particular, and a general normative judgment or principle.

This concludes my defence of cognitive universalism. In the next and final section, I briefly discuss its wider implications.

## 5 Philosophical implications

As far as I can tell, cognitive universalism is compatible with a wide range of metaethical views. It is of course a form of cognitivism, and just like Moore, it takes the basic normative concepts to resist definition or analysis in descriptive terms.<sup>13</sup> So it is at odds with non-cognitivism and analytical naturalism. But we need not follow Moore in inferring that normative *properties* also resist reduction in terms of descriptive properties. So, for example, it should be possible for a *culture relativist* to accept that normative claims about kinds are basic, in the way I have described, and add that what it is for a kind to have a normative property, relative to a given culture, is for there to be a relevant norm about the kind in force in that culture. Or a *constructivist* could say that what it is for a kind to have a normative property, for some subject, is for the subject to have a relevant evaluative attitude

---

<sup>12</sup> This point is from Fine (2005).

<sup>13</sup> I am ignoring the question of whether, as Moore thought, some of the normative concepts can be defined in terms of the others.

towards the kind. And so on. These do not seem like big departures from what defenders of such views believe already.

I want to focus on the implications for nonnatural realism, which are rather interesting. I said in the introduction that this view is often thought to have trouble explaining normative-descriptive supervenience. And I said that “explaining supervenience” is really two tasks: i) explaining why normative properties supervene on descriptive properties, and ii) explaining why this pattern is analytic.

Cognitive universalism takes care of the second task. It explains why, as a conceptual matter, any referent of “good”, “bad”, “right” etc, must supervene on descriptive properties. But what about the first task? We now know that, unless nonnatural values and reasons supervene on descriptive properties, they will be disqualified as referents of our normative predicates. But this does nothing to *explain why* they supervene on descriptive properties; it just reinforces the point that they’d better do it.

Another way to make the same point is this: we now know why it is conceptually necessary that value supervenes on descriptive properties. Since conceptual necessity is stronger than metaphysical necessity, it follows that it is metaphysically necessary that value supervenes on descriptive properties. But it does not follow that *nonnatural normative properties* supervene on descriptive properties with metaphysical necessity. That only follows if we assume that value = such-and-such nonnatural property. But the first supervenience worry is precisely a *challenge* to that identity claim. The challenge is that, *unlike* value, it is not clear that nonnatural normative properties would be well-behaved, supervenience-wise. So in order to meet the challenge, we want an independent explanation of why they are well-behaved.

Furthermore, there is *Hume’s Dictum* to worry about: the principle that there be no “necessary connections between distinct existences”. It is not obvious exactly what that means (c.f. Wilson (2010)), but in the present context it boils down to a ban on distinct, cointensional (necessarily coextensive) properties. In other words: if, *in every possible world*, all the Fs are Gs and vice versa, then F and G are the same property. For if F and G are really different properties, surely it is *possible* for something to be F without being G or vice versa?

Nonnatural normative properties violate this principle because, given supervenience, they are each going to have a cointensional descriptive property. Take for example goodness: just look at all the good things in all the possible worlds, and put each of their profiles of descriptive properties and relations together in a long, disjunctive property D. Given supervenience, goodness and D are cointensional: every possible good thing is also D, and every possible D thing is also good.<sup>14</sup>

---

<sup>14</sup> See Jackson (1998, ch. 5) for a rigorous statement of this point.

The worry, then, is that we have no good explanation of why nonnatural values or reasons would supervene on descriptive properties, and that, supposing they *do* supervene, they will violate (a version of) Hume's Dictum, by having cointensional descriptive properties (from which they are nevertheless supposed to be distinct). So if nonnatural realism is to be plausible, it needs to explain why non-natural values or reasons supervene on descriptive properties, and explain it in a way that either gets around Hume's Dictum, or else makes it plausible that (the relevant version of) the Dictum is false.

The crucial thing to notice about this challenge is that it is stated in terms of the normative properties that *particular* things have. What we want explained is why the values of *particular things* supervene on their descriptive properties; and it is the *particular-applying* normative properties that will have cointensional descriptive properties.

Many people seem to think the nonnatural, irreducibly normative properties that Moore, Parfit, Nagel, etc, posit take particular things as their bearers. For example, that Moore posited a nonnatural property ("goodness") that is instantiated by things like *Bill's pleasure at time t from looking at picture x*, or *Bob's knowing at t that p*. Or that the irreducibly normative reasons Parfit posits are relations between *particular* facts and responses available to particular agents. For example, that the fact *that Bill is in pain* is a reason *for Bob to give him an Aspirin*.

If that is the view, then it will indeed be mysterious why these properties should supervene on descriptive properties. Suppose we have, as a basic, irreducible fact, that the fact that *Bill* is in pain (and *Bob* is around with a spare Aspirin), is a reason for *Bob* to volunteer the Aspirin. Then why would it be impossible for Susan and Tracy, say, to find themselves in a similar predicament, but without Tracy having a similar reason? (Supposing, again, that the reason-givingness of *Bill's* pain is a *basic* fact, not deriving from some underlying, general normative fact.)

But I can see no reason why the view should take that form. The main intellectual motivation behind nonnatural realism has always been a strong commitment to certain pre-theoretical, first-order normative beliefs. And now that cognitive universalism has shown us that these beliefs bottom out in truth conditions about what kinds of things have value or provide reasons, it is only natural to direct one's metaphysical commitment accordingly: to facts about what kinds of things have value or provide reasons.

And if we take up that metaphysical commitment, then there is no need to posit, *in addition*, irreducible normative properties that take particulars as their bearers. Instead, we can give a *reductive* account of the particular-applying normative properties. For example, we can say that the property *goodness* (the one that takes particulars as bearers) is just the property *being a token of a good kind*.

On this reductive view, we *can* explain why the values of particular things supervene on their descriptive properties. Whether or not a particular thing is good

will depend on two things: i) which (descriptive) kinds it tokens, and ii) which (descriptive) kinds are good. The first part supervenes on the thing's descriptive properties in a real and obvious way: it is a token of a given kind just in case it instantiates the corresponding property. The second part, concerning which kinds of things are good, also supervenes on descriptive properties, but in a trivial, uninteresting sense. Facts about what kinds of things are good, nonnatural realists say, are *necessary*, and necessary facts trivially supervene on everything (there can be no difference in the necessary facts without a difference in the Y-facts, for any Y you like, because there can be no difference in the necessary facts, period).

This view will, to be sure, leave *something* unexplained, and it will contain *some* necessary connections between distinct properties. But it is crucial to see exactly what is left unexplained, and exactly what these necessary connections are like. What is left unexplained is not supervenience, but rather (some of the) facts about which kinds of things are good, bad, wrong, etc. For example, that causing needless pain is wrong; that happiness is good; that suffering is bad. Some of these facts may be explainable in terms of the others, but some of them are going to be basic, and admit of no further explanation.

Everyone agrees it is a desideratum on metaethical theories that they should explain why the values of particular things supervene on their descriptive properties. Nonnatural realism can do that, using the reductive account of particular-applying normative properties, and appealing to facts about the values of kinds. So the question is whether it is OK to leave these latter facts unexplained. In other words, is it also a desideratum on metaethical theories that they should *explain why* suffering is bad, or why happiness is good, and so on? To me these seem like very good places for explanations to end. But it is hard to argue about where explanations should end; so let us just record that the realist view developed here will have such commitments.

But what about *the necessity* of these facts? Will they not give us “necessary connections between distinct existences”? Yes they will, but it is important to see exactly what these connections are like. Actually they come in two forms. First, particular-applying normative properties will have cointensional descriptive properties, like goodness and D, as discussed above. But (particular-applying) goodness, recall, is not a fundamental property; it is just the property *being a token of a good kind*. That this property necessarily co-occurs with D is just a trivial consequence of its definition, given that facts about the values of kinds are necessary.

So the interesting necessary connections are these latter facts themselves. Take agony and badness, for example. The necessary connection here is not *co-occurrence*, but *instantiation*. It is not that agony and (kind-applying) badness are instantiated by the same things; it is that agony itself instantiates badness, that agony *is* bad. In other words, the necessary connection is that a first-order universal has a second-order universal.

But necessary connections between first- and second-order universals are legion. Crimson, for example, has the property of being a shade of red. And it has that property in every world.<sup>15</sup> That is not strange at all; of course we need not check, with any given world, to see whether crimson is a shade of red there. Likewise, we need not check, with any given world, to see whether agony is bad there. That agony is bad has to do with how awful it is to be in agony. But what it is like to be in agony does not change from world to world. We need not check to see whether it is horrible to be in agony in a world; for if it were not horrible, it would not be agony. Likewise, we need not check with a given world to see whether *treating another as a mere means*, say, is wrong there. If acting this way is wrong, that is a fact about this interpersonal relationship, which is the same in every world. In general, it is plausible that the basic normative facts are invariant from world to world, because they are facts about things that do not vary from world to world.

I want to end, however, by mentioning a variation on the view that may be even more plausible. The view just sketched requires metaphysical commitment not only to nonnatural normative properties, but also to *kinds* to serve as their bearers. If we have kinds in our metaphysics anyway, that is of course no problem, but not everybody does. However, suppose we have *properties* in our metaphysics, for independent reasons.<sup>16</sup> Then we can suit the view to our liking, as follows: instead of saying the metaphysically basic fact that makes “causing needless pain is wrong” true is that the kind *causing needless pain* has the property *wrong*, we can say it is that the property *causes needless pain* has the property *wrongmaking*. Rather than posit the kind *pleasure* to instantiate the property *good*, we can posit the property *pleasant* to instantiate the property *goodmaking*. And so on. For each basic, normative concept *F* that applies to kinds, we posit a corresponding *F-making* property that applies to properties.

On this second view, the property *goodness* (the one that takes particulars as bearers) will just be the property *having a goodmaking property*. The property *wrongness* will

---

<sup>15</sup> On an Aristotelian view, on which universals only exist in worlds where they have instances, all we could say would be that *crimson* has *being a shade of red* in every world in which it exists. But then again, on this view, *agony* likewise has *badness* only in all the worlds *agony* exists in, so the case is parallel.

Another possible objection is that *being a shade of red* applies to other things besides crimson, so we do not have the one-one necessary coexistence that strike some as particularly worrisome. But then again, other things besides agony are bad, so the case is again parallel.

<sup>16</sup> I do not mean to suggest that there is a deep metaphysical difference between properties and kinds. There seems to be a shallow difference, indicated by the linguistic phenomena below (for example that we can say “is wrong” about the kind *murder*, but have to say “is wrongmaking” about the property *murderhood*). I am not sure what to think about this.

just be the property *having a wrongmaking property*. And so on. The explanation of supervenience proceeds as before.

The possible attraction of this second view is that it replaces metaphysical commitment to kinds with commitment to properties. But either way, these views need *universals*, of one flavour or another, to serve as bearers of the basic normative properties. If the suggestions I have made in this last section are on the right track, then, the lesson the nonnatural realist should draw from cognitive universalism, is that she should be a platonist.

## References:

- Blackburn, Simon (1971); "Moral Realism" in *Morality and Moral Reasoning*, ed. John Casey, London: Methuen.
- Blackburn, Simon (1985); "Supervenience Revisited" in *Exercises in Analysis: Essays by Students of Casimir Lewy*, ed. Ian Hacking, Cambridge, Cambridge University Press
- Blackburn, Simon (1993); *Essays in Quasi-Realism*, OUP
- Carlson, Gregory N (1977); *Reference to Kinds in English*, PHD Dissertation, University of Massachusetts, Amherst
- Davidson, Donald (1980); *Essays on Actions and Events*, OUP
- Drai, Dalia (2000); "Moral Supervenience and Moral Thinking" in *Disputatio*, vol. 8
- Dreier, James (1992); "The supervenience argument against moral realism" in *The Southern Journal of Philosophy*, Vol. XXX, No. 3
- Fine, Kit (2005); *Modality and Tense*, OUP
- Jackson, Frank (1998); *From Metaphysics to Ethics*, OUP
- Gibbard, Allan (1990); *Wise Choices, Apt Feelings*, Harvard University press
- Gibbard, Allan (2003); *Thinking How to Live*, Harvard University Press
- Leslie, Sarah-Jane (2008); "Generics: Cognition and Acquisition" in *The Philosophical Review*
- Lewis, David (1983); "New Work for a Theory of Universals" in *Australasian Journal of Philosophy*, vol. 61, 4
- Oddie, Graham (2005); *Value, Reality, Desire*, OUP
- Hare, R. M. (1952); *The Language of Morals*, OUP
- Korsgaard, Christine (1996); *The Sources of Normativity*, Cambridge University Press
- Koslicki, Kathrin (1999); "Genericity and Logical Form" in *Mind and Language*, Vol.14, No.4
- Kramer, Matthew (2009); *Moral Realism as a Moral Doctrine*, Wiley-Blackwell
- Liebesman, David (2011); "Simple Generics" in *Nous*, vol. 45:3
- Miller, Alexander (2003); *A Contemporary Introduction to Metaethics*, Polity Press,

- Cambridge
- Moore, G. E. (1903); *Principia Ethica*, Cambridge University Press
- Moore, G. E. (1922); “The Conception of Intrinsic Value” in *International Library of Philosophy: Philosophical Studies*, Routledge and Kegan Paul
- Nagel, Thomas (1986): *The View from Nowhere*, OUP
- Nagel, Thomas (1996): “Universality and the reflective self” in Korsgaard, Christine: *The Sources of Normativity*, Cambridge University Press
- Parfit, Derek (2011); *On What Matters*, OUP
- Smith, Michael (1994); *The Moral Problem*, OUP
- The Stanford Encyclopedia of Philosophy (2008); *Moral Non-Naturalism*. Available (as of January 10th 2012) at: <http://plato.stanford.edu/entries/moral-non-naturalism/>
- Wilson, Jessica (2010); “What is Hume’s Dictum, and Why Believe It?” in *Philosophy and Phenomenological Research*, vol. 80, p. 595-637