

How to Be a Moral Platonist

Knut Olav Skarsaune

This chapter will develop a form of non-natural realism about value, arrived at by way of reflection on the so-called “supervenience objection” to the view. Though I will argue the objection fails as an *objection*, I think it is a first-rate *challenge*; one that forces us to come back with a clearer and hopefully better statement of our position.

The supervenience objection has been discussed and developed so widely that it is probably quixotic to try to give a statement that will satisfy everyone. One reason to hope many will find their favourite version addressed is that I take up, not one, but two objections. The bulk of the chapter responds to Simon Blackburn’s (1971, 1985) challenge, which is to explain *why it is analytic* that normative properties supervene on descriptive properties. In the last section, I take up a more recent, metaphysical challenge, which asks how normative properties can be *fundamental* and at the same time supervene on other properties. Does not that violate *Hume’s Dictum*, the ban on “necessary connections between distinct existences”?

The heart of my response to Blackburn is an idea I got from Kit Fine.¹ It is that normative predicates express subtly different senses when they are applied to *particular* things and to *kinds*, respectively. The kind-applying senses are basic and primitive, while the particular-applying senses are defined in terms of them. A particular thing is good_{PAR}, for example, just in case it is a token of a good_{KIN} kind. But then, since two descriptive twins are tokens of exactly the same kinds, they must either both be good_{PAR} or both not be. Supervenience falls out of the definitions of the particular-applying normative concepts.

This account is a cognitivist analogue of R. M. Hare (1952) and Allan Gibbard’s (1990) non-cognitivist accounts. 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

¹ In discussion. Many thanks to Fine for the idea; all mistakes, of course, are mine.

Gibbard, endorsement of a *norm*). My account applies the same 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 chapter. 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.

In the last part of the chapter I turn to metaphysics. I argue that non-natural realists should just repeat, at the level of properties, the explanation of supervenience that cognitive universalism provides at the level of concepts. Just as they should say that the basic normative concepts apply to kinds, so they should say that the basic normative properties are second-order; they do not apply to particular things, but rather to kinds, or perhaps to properties. I argue, in other words, that non-natural realists should be platonists.

10.1 BLACKBURN'S CHALLENGE

Blackburn's challenge is to explain why normative judgments exhibit a peculiar combination of independence and dependence vis-à-vis descriptive judgments. Let us start with the independence. From G. E. Moore (1903) onwards, non-natural realists have said that at least some normative concepts are primitive, undefinable, basic. Perhaps some of them can be defined in terms of others, perhaps they form an interdefinable cluster, but no normative concept can be defined in entirely non-normative terms. Prominent candidates for primitiveness include GOOD and BAD (Moore); BEING A REASON FOR (Parfit, Nagel, Scanlon); and OUGHT (Broome, Wedgwood). I want to be neutral on this issue, and will switch liberally between different normative concepts in my exposition.

The independence at issue is more than primitiveness. RED, for example, could be primitive and yet not have the relevant kind of independence from descriptive concepts, being one of them itself. But non-natural realists further maintain that normative concepts are *different in kind* from descriptive concepts. Descriptive concepts are used to judge how things are; normative concepts are used to judge how they should be; and these are different kinds of judgment.

The upshot is that descriptions never entail evaluations. There is no description D such that

- (1) *a* is D.

entails

(2) *a* is good.

Or as Blackburn puts it: ‘There is no moral proposition whose truth is entailed by any proposition ascribing naturalistic properties to its subject’ (1993: 116).

So far independence. But there is also a close connection between normative and descriptive judgment. As far as I know, every prominent non-natural realist accepts that it is in some sense a conceptual truth that things have their value *in virtue of* being the way they are descriptively, and that value therefore supervenes on descriptive properties.²

Speaking now in my own voice, I believe the best way to make this precise is as follows. What I want to say is not that any grand supervenience principle is itself analytic, but rather that ordinary English sentences and inferences that exemplify certain principles are analytic.³ Keeping in mind throughout that descriptive likeness includes both properties and relations, I take it that the inference from

(2) *a* is good.

(3) *b* is descriptively exactly like *a*

to

(4) *b* is good.

is analytic, and likewise

(5) If Ted is good, then it is impossible to be just like Ted in every descriptive respect and not be good.

Sentence (5) exemplifies

² Notice that Blackburn uses “naturalistic” to describe the supervenience base. While I have no problem with this, it tends to invite needless distractions having to do with divine command theory. I therefore use “descriptive” instead. In order to head off Sturgeon’s (2008) worries, let me stress that I am not referring to the properties picked out by this or that descriptive vocabulary. “Descriptive property” is meant to convey a metaphysical notion, probably primitive, but sometimes elucidated with the phrase “ways things can be”.

³ A sentence “exemplifies” a principle just in case its negation contradicts the principle; an inference exemplifies it just in case the premises together with the negation of the conclusion contradict it. Of course, in this sense the supervenience principles exemplify themselves, but as we shall see, they are not “ordinary English sentences”, hence they fall outside the analyticity claim.

Strong normative-descriptive supervenience: for all possible x, y : if x and y are alike in every descriptive respect, then x and y are alike in every normative respect.

while the inference from (2) and (3) to (4) exemplifies both strong and

Weak normative-descriptive supervenience: for all x, y in the same possible world: if x and y are alike in every descriptive respect, then x and y are alike in every normative respect.

Strictly speaking, the principles themselves are not analytic, because they are not true in virtue of the meanings of the words that figure *in them* alone; perhaps they are in some sense true in virtue of the meanings of “good”, “bad”, etc., but these words do not figure in the principles. There, rather, we find the term of art “normative”, and I doubt that *its* meaning guarantees supervenience. But speaking loosely, I will say that both weak and strong supervenience are analytic, in the derivative sense that ordinary English sentences and inferences that exemplify them are analytic.

Are they analytic in this sense? 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 logical and not merely a moral mistake that had been made. (1993: 116)

R. M. Hare:

we cannot say ‘P is exactly like Q in all respects save this one, that P is a good picture and Q not’. If we were to say this, we should invite the comment, ‘But how can one be good and the other not, if they are exactly alike? There must be some *further* difference between them to make one good and the other not.’ Unless we at least admit the relevance of the question ‘What makes one good and the other not?’ we are bound to puzzle our hearers; they will think that something has gone wrong with our use of the word ‘good’. Sometimes we cannot specify just what it is that makes one good and the other not; but there always must be something. Suppose that in the attempt to explain our meaning we said: ‘I didn’t say that there *was* any other difference between them; there is just this one difference, that one is good and the other not. Surely you would understand me if I said that one was *signed* and the other not, but that there was otherwise no difference? So why shouldn’t I say that one was *good* and the other not, but that there was otherwise no difference?’ The answer to this protest is that the word ‘good’ is not like the word ‘signed’; there is a difference in their logic. (1952: 81)

I will not try to defend these analyticities, or analyticity in general, in this chapter. I will only report that I share Hare and Blackburn’s intuitions. If

an interlocutor repeatedly flouted supervenience, by calling some things “good” and other things “not good” on the basis of no descriptive difference whatever, then I would be left with nothing else to think than that her 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”.

The dialectical situation is this: I will shortly present Blackburn’s objection, namely that non-natural realists *cannot explain* these analyticities, and then respond to his objection by offering an explanation. It is somewhat ironic that, if I should succeed perfectly, my reward will be to get in trouble with those who deny the analyticities in the first place, perhaps because they reject the notion of analyticity altogether. But *c’est la vie*; one has to start theorizing from some conception of the data, as best one can discern them.

The stage is set for Blackburn’s objection:

if A has some naturalistic properties, and is also good, but its goodness is a distinct further fact not following from the naturalistic features, and if B has those features as well, 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. (1993: 118–19)

The objection is that non-natural realists have no good explanation of the analyticity of supervenience, exemplified by entailments like the one from (2) and (3) to (4).⁴ The view that normative predicates are primitive or unanalysable seems to prevent us from pointing to any feature of their

⁴ Blackburn later (1985) restated the argument as follows: Suppose we judge a thing to be good. We are then committed to there being some collection of descriptive properties and relations that underlie its goodness. Put all these together in a big, conjunctive property F . Include in F also all its normatively relevant *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. If descriptions never entail evaluations, there will be *conceptually possible worlds* in which the F ’s are not good, as well as worlds in which the F ’s are good. But since *every* normatively relevant property and relation, both positive and negative, is included in F , there will not be 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, then 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 the original version, so I will not discuss it any further in the text.

meaning that would produce this entailment. The only option seems to be to take supervenience as a further primitive; to say that it is just a basic fact about the meanings of normative terms that they behave this way. That is what Allan Gibbard, for example, thinks the non-natural realist has to say:

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: 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: 184)

I agree with Blackburn and Gibbard that this would be a weak position. It is not an inviting view to maintain that there are just three things to say about, say, the concept GOOD, namely (a) it is primitive, (b) it is not descriptive, and (c) it cannot apply to one but not the other of two descriptive twins. If the concept is primitive, where does this restriction come from?

However, I believe (c) has an underlying explanation, to which I now turn.

10.2 COGNITIVE UNIVERSALISM

I begin with weak supervenience, the rule about evaluation of worldmates. I will add an hypothesis in section 10.5 that will take us from the weak to the strong principle.

The account begins by noticing that normative predicates can grammatically apply both to *particulars* (dated, non-repeatable things) and to *kinds* (timeless things that can have instances). We can say that Florence Nightingale was good or that altruism is good; we can say that the Iraq war was wrong or that larceny is wrong. One might think this is merely a surface phenomenon, because normative claims about kinds are covertly quantificational. For example, “Altruism is good” might be analysed as saying that all, or typical, or that in general, *instances* of altruism are good. On such analyses, notice, some quantifier binds a variable, the variable ranges over particulars, and the predicate is applied to the variable. After the sentence has been interpreted, then, the normative predicate is applied to particulars also in general claims.

I will defend a different account of such claims, on which the normative predicates are genuinely applied to kinds. For example, “Altruism is good”

does not say that all/typical/in general instances of altruism are good. It says that *altruism*, a kind of motivation, is itself good.

Before I explain this proposal, I should mention that we are not compelled to treat all such claims the same way. It could be that some are quantificational and others kind-referring, or that some allow both readings. It seems plausible that many *non-normative* uses of “good” or “bad” should be analysed as quantificational generic claims or “generics”. For example, perhaps “Apple pie is good” should be analysed: in general, if x is an instance of apple pie, then x is good. Perhaps some apparently kind-referring normative claims are also quantificational generics, or at least permit such readings. The account I will propose is perfectly compatible with this possibility, albeit in a roundabout way.⁵ But for ease of exposition I will present a “clean” view on which the relevant claims are always kind-referring.

The proposal, then, is that normative predicates can apply both to particulars and to kinds, not just superficially but also after the sentence has been interpreted. The next question is how to explain the logical relationship between claims about kinds and claims about particulars. This issue may take a moment to register, given how used we are to quantificational analyses of general claims. Explaining logical relations between general and particular claims is just what quantifiers do. For example, taking us from “All F’s are G” together with “ a is F” to “ a is G” is just what “all” does. But suppose “Murder is wrong” has the simple logical form $\text{WRONG}(\textit{murder})$. Suppose further that “ a is an instance of murder” and “ a is wrong” have the logical forms $\text{INSTANCE}(a, \textit{murder})$ and $\text{WRONG}(a)$, respectively. In that case we lose the logical relation between these claims; the three formulas in small caps are logically independent of each other.

This is an appropriate result in some other cases. Claims of the form “Kind K is F” do not always imply anything about whether individual K ’s are F. For example, the Monsanto Company has patented the genetically modified corn *Genuity VT Triple PRO*. Which seems to make

(6) Monsanto owns Genuity VT Triple PRO.

true. But nothing follows about who owns the kind’s instances; (6) is compatible with Monsanto being sold out of Genuity ..., all the physical corn thus being the property of others.

⁵ Once I distinguish the *pro tanto* and all-things-considered readings (discussed below), I am personally not able to get an exception-allowing, generic reading of “Lying is wrong”, but suppose we want one. On my account it would be $\text{GEN}_X [\text{LYING}(x)]$ ($\text{WRONG}_{\text{PAR}}(x)$), which, given the definition that will shortly be offered in the text, would in turn be analyzed as $\text{GEN}_X [\text{LYING}(x)] (\exists K [\text{TOKEN}(x, K) \ \& \ \text{WRONG}_{\text{KIN}}(K)])$. “GEN” is a quantifier that means, roughly, “in general”.

But in the normative case, there are clearly logical relations between claims about kinds and claims about particulars. It would be incoherent to say that murder is wrong but that each murder has no tendency to be wrong. If this is not explained by a silent quantifier in “Murder is wrong”, then how is it explained?

A natural idea would be to reinstate the quantificational approach at the level of lexical semantics, in the following way: to say that “wrong”, *when applied to a kind*, does not express the concept WRONG, but the concept WRONG*, which in turn has the definition BEING SUCH THAT ALL ONE’S INSTANCES ARE WRONG. In other words, the *sentence* would not have a quantifier in it, its logical form would just be WRONG*(murder). But the definition of WRONG* would in turn have a quantifier in it, which generates the entailment from “Murder is wrong” and “*a* is a murder” to “*a* is wrong”. Or perhaps one would prefer to define WRONG* as BEING SUCH THAT, IN GENERAL, ONE’S INSTANCES ARE WRONG, in which case “Murder is wrong” together with “*a* is a murder” would provide some kind of defeasible support for “*a* is wrong”.

However, for reasons I will give later, I believe we should instead explain the logical relations *from the opposite direction*; we should take normative predicates to express simple senses when they are applied to kinds, and complex senses when they are applied to particulars.

Concerning the kind-applying senses, in other words, I suggest that non-natural realists should carry on saying the things we have always said about normative concepts (primitive, irreducible, simple, basic, fundamental). But for the particular-applying senses, we should not say those things. These senses are, on the contrary, complex and definable. Not, though, in terms of non-normative concepts, but in terms of the kind-applying ones. For example, let GOOD_{PAR} and GOOD_{KIN} be the particular- and kind-applying senses, respectively, of “good”. Then we can define the former in terms of the latter, as follows:

$$\text{CU: } \text{GOOD}_{\text{PAR}}(x) \longleftrightarrow_{\text{def}} \exists K [\text{TOKEN}(x, K) \ \& \ \text{GOOD}_{\text{KIN}}(K)]$$

A particular is good_{PAR} just in case it is a token of a good_{KIN} kind.

The variable “*K*” ranges over descriptive kinds. *War*, for example, is an event kind; *eating bananas* is an act kind, *being happy* is a kind of mental state. Importantly, there is no restriction on how general or specific the kinds are. The range of *K* includes *eating bananas while sitting on a train passing by a lake*, and even kinds that cannot be expressed in English at any length.

However, we should not include so-called *haecceitic* kinds in the range of *K*. A kind is haecceitic if it concerns a specified individual. 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. In other words, we should understand “alike in every descriptive respect” in our supervenience principles to mean “alike with respect to all qualitative descriptive properties and relations”. My impression is that that is what Hare, Blackburn, and most participants in the subsequent debate have intended.⁶

Like properties, kinds can have instances/tokens. 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_{KIN} applies to a kind, then GOOD_{PAR} applies to *every* instance of the kind. That might seem too strong. For example, we might be inclined to accept both

(7) Knowledge is good.

and

(8) Oedipus’s 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 widely agreed that normative predicates can express both *pro tanto* and *all-things-considered* senses.⁷ And we

⁶ One exception is Matthew Kramer (2009: ch. 10). He points out that many religious believers think *being pleasing to God* is normatively relevant, but they can hardly be accused of conceptual confusion. He also imagines a man who favours 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 religious believers think God’s opinion matters because he has certain qualitative properties (power, wisdom) and stands in certain qualitative relations to them (has created them, cares about them). It is not clear that anyone thinks God’s haecceitic identity makes a difference, so that, for example, we could have another qualitatively identical world, with a qualitatively identical creator, but the people over there have no reason to obey *their* creator.

The thought-experiment with the Francophile is also weak. We need to imagine two qualitatively identical countries, France and Schmance, and our man must stand 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 acts that benefit France “good” and acts that benefit Schmance “not good”. Then I would simply repeat Hare’s intuition from section 10.1: I would be left with nothing else to think than that “something has gone wrong with his use of the word ‘good’”.

⁷ 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 each of them.

must distinguish these *before* we apply CU. For example, in (7) we have PRO-TANTO-GOOD_{KIN}, and in (8) we have ALL-THINGS-CONSIDERED-GOOD_{PAR}. Under those readings, (7) and (8) are jointly compatible with CU.

This is not an ad hoc move: the distinction between *pro tanto* and all-things-considered senses 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 (8) is true, then you will not get (7) to be true if you force the all-things-considered reading. And if you think knowledge is *pro tanto* good, you will not get (8) to be true if you force the *pro tanto* reading.

So far I have discussed “good”, but I think parallel analyses apply at least for “bad”, “right”, “wrong”, “just”, “unjust”, “fair”, and “unfair”, or more accurately for both the *pro tanto* and the all-things-considered senses of these.⁸ CU, then, is a general recipe for defining particular-applying senses of normative predicates in terms of their kind-applying senses.

CU has the same structure as 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 commitment is existentially quantified.⁹ 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, I will call it *cognitive universalism*.

The view now presented, I will proceed to give some reasons to accept it. The first is that it explains the analyticity of weak supervenience, and as we shall see in section 10.5, also that of strong supervenience, when we combine it with a further hypothesis about normative predicates. But let us start with the weak. The challenge is to explain things like the analytical entailment from

- (2) *a* is good.
- (3) *b* is descriptively exactly like *a*.

to

- (4) *b* is good.

⁸ I think analyses similar in spirit, but perhaps different in detail, apply for “reason” and “ought”. And I suspect that aesthetic predicates like “beautiful” behave in this way too.

⁹ 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).

The explanation is this. From (2), 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_{KIN} . From (3) 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_{KIN} , and so by UC,

(4) b is good.

It should be straightforward to see how this generalizes 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).

We have, in effect, 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_{KIN} and a defined concept GOOD_{PAR} . Weak supervenience just falls out of the definition of GOOD_{PAR} .

10.3 BEDFELLOWS

It might seem extravagant to posit a systematic duplication of predicate senses, just in order to explain the analyticity of supervenience. But in this section, I will present independent linguistic evidence for this duplication. I believe we should posit it anyway; not just for normative predicates but indeed across language.

The issue is tied to another feature of the account just given; its division of labour between sentence semantics and lexical semantics. The work of explaining truth-conditions and inferential relations between sentences is not done in a single step, at the level of logical form; a significant part of the job is relegated to lexical semantics. I will begin by saying something more to motivate and explain this general approach.

It is nearly uncontroversial that *kind-selecting* predicates like “extinct”, “widespread”, “invented by”, and “legalize” only take kinds as subjects. Consider

(9) Skateboarding was invented by bored surfers.

(10) Cain invented fratricide.

The predicates are true of the act kinds in question, 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 we can come up with quantificational truth-conditions for

(11) Mehmed II legalized fratricide.

(What makes it difficult is that (11) 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 produce 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 (11) a rather simple logical form, and instead assign “legalize” a rich lexical semantics.

This point is reinforced by the fact that the inference from (10) and (11) to

(12) Mehmed II legalized 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*, NP's occurring together with kind-selecting predicates really do refer to kinds. For example, the logical forms of (11) and (12) are just

(11*) LEGALIZED(*Mehmed*, *fratricide*)

(12*) $\exists x$ [LEGALIZED(*Mehmed*, *x*) & INVENTED(*Cain*, *x*)]

This view does not imply that (11) is true because Mehmed II performed some mysterious act on an abstract object (the act kind *fratricide*). Clearly, what *ultimately* makes (11) true is that he performed certain particular acts; perhaps he signed a document or made an announcement in the town square. Similarly, (9) is true because a gang of bored surfers performed some particular, pioneering acts of skateboarding. That is how you legalize, 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*.¹⁰ So for example, because we understand what “widespread” means, we understand that the worldly truth-condition of “Bedbugs are widespread in Brooklyn” does not involve some abstract object being spread out over Brooklyn, but rather that there be concrete bedbugs at many locations across the borough.

Once we employ this division of labour between sentence and lexical semantics, and make the corresponding distinction between logical form and worldly truth-conditions, it becomes extremely plausible that ordinary particular-applying predicates, like “have four legs” or “fifty feet tall” are also capable of expressing kind-applying senses. Consider

(13) *Panthera tigris* has four legs.

(14) There was a dinosaur species that was fifty feet tall.

It is extremely plausible that “has four legs” in (13) is predicated of the kind *Panthera tigris*, and that “fifty feet tall” in (14) is predicated of a variable that the rest of the sentence relates to a dinosaur species.

But (13) does not say that tiger-kind has four legs in the same sense as Bo the dog. Roughly, what (13) says about tiger-kind is that *its genetically normal, unmolested tokens have four legs*. Likewise, what (14) says about the dinosaur species is that *its healthy adult tokens were about fifty feet tall*. Someone who does not understand this does not understand the claims being made. But since this is not encoded in logical form, the work must be done by the predicates instead.

The view, then, is that the logical forms (ignoring time) are

(13*) HAVE-FOUR-LEGS*(*Panthera tigris*)

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

The division of labour I discuss here, between sentence 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.

(14*) $\exists x [\text{DINOSAUR SPECIES}(x) \ \& \ \text{FIFTY-FEET-TALL}^*(x)]$

where HAVE-FOUR-LEGS* is the concept BEING SUCH THAT ONE'S GENETICALLY NORMAL, UNMOLESTED TOKENS HAVE FOUR LEGS, and FIFTY-FEET-TALL* is the concept BEING SUCH THAT ONE'S HEALTHY ADULT TOKENS ARE ABOUT FIFTY FEET TALL.

The view is not that all these predicate senses are lexicalized. Plausibly, the usual case is that the particular-applying senses are lexicalized, while suitable kind-applying senses are constructed ad hoc, in order to get reasonable interpretations. For example, if I say "The English Setter weighs between 40 and 50 pounds", you construct a suitable kind-applying sense of "weigh between 40 and 50 pounds" on the spot, in order to get a sensible interpretation.

This account predicts that, if we apply a predicate to a particular-referring and a kind-referring NP at the same time, the result should be at least mildly zeugmatic. Zeugmas are not ungrammatical, but feel a bit like puns, in that a word is used to mean two different things at the same time. Since the particular-applying and kind-applying senses are so closely related, the effect is not likely to be as strong as in "He took his hat and his leave", but it should be perceptible. As indeed

(15) Bo and *Panthera tigris* both have four legs.

(16) *Equus zebra* and uncle Bob live in Africa.

testify.

I have given some reasons to think it is a common phenomenon that predicates express different senses when they are applied to particulars and to kinds, respectively. And I have explained the general framework I think this kind of view should be located in. Turning now to normative language in particular, I will present some evidence that normative predicates can be applied to kind-referring NP's (I assume it is clear that they can be applied to particular-referring NP's, as in "Florence Nightingale was good").

Since "widespread" is a kind-selecting predicate, "murder" refers to a kind in

(17) Murder is widespread.

But then it seems that "wrong" is applied to a kind in

(18) Murder is widespread and wrong.

(19) Murder is widespread even though it is wrong.

Likewise, we can apply "good" to a variable that the rest of the sentence relates to a kind:

(20) Some kinds of charity are good.

There is independent evidence, then, that normative predicates can be applied both to particulars and to kinds, not just superficially but in logical form. The main reason to think that they express *different* senses in the two uses has already been given: otherwise it would be hard to explain the logical relations between the two sorts of claim. In addition, and finally, I suggest there is a feeling of zeugma when we apply a normative predicate to kind-referring and particular-referring NP's at the same time:

(21) The war in Iraq and sexual harassment are wrong.

(22) Philanthropy and Warren Buffett are virtuous.

10.4 WHY THE KIND-APPLYING SENSES 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 am neutral on this 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.

This is because the kind-applying senses of normative predicates are basic. So a normative claim about a kind, like “Murder is wrong”, will not be reduced in terms of particulars at any level of semantic analysis. Its logical form is $\text{WRONG}_{\text{KIN}}(\textit{murder})$ and its worldly truth-condition is also $\text{WRONG}_{\text{KIN}}(\textit{murder})$.

On the other hand, applications of normative concepts to particulars are analysed away at the level of lexical semantics. However, they are given

mixed worldly truth-conditions, in part descriptive, in part normative. For example, “Bob is good” has the logical form $\text{GOOD}_{\text{PAR}}(\text{Bob})$. Applying the definition of GOOD_{PAR} , we get the worldly truth-condition

$$\exists K(\text{token}(\text{Bob}, K) \ \& \ \text{good}_{\text{KIN}}(K)).$$

Bob is a token of some descriptive kind or other, such that that kind is good_{KIN} . So the *descriptive* part of the truth-condition is still about the particular, Bob. But the *normative* part, saying that the relevant kind is good_{KIN} , 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 have already given the main reason to believe this; it explains the analyticity of (weak) supervenience. 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. Moral 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.

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 is 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 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 generalizations. But at the bottom of all this epistemic activity is the single ur-judgment that an act is right if and only if it maximizes 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 moral 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 Srebrenica Massacre was wrong. But this fixed point is not really that a particular act *de re*, the Srebrenica Massacre, was wrong. For imagine we became convinced by some internet conspiracy theory that the “massacre” is in fact a fabrication of Western propaganda. 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 eight 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 Srebrenica 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 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.¹¹

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 should expect if cognitive universalism is true. Since normative claims about kinds are not generalizations over their instances, they are not justified in the way generalizations 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.

10.5 WHY STRONG SUPERVENIENCE IS ANALYTIC

Earlier, I postponed discussion of strong supervenience and its analyticity. My formulation of the principle was:

Strong normative-descriptive supervenience: for all possible x, y : if x and y are alike in every descriptive respect, then x and y are alike in every normative respect.

What I take to be analytic, again, are not such principles themselves, but ordinary claims and inferences that exemplify them. Since strong supervenience coordinates the normative properties of individuals in different worlds, we have to look to modal or counterfactual claims in order to find sentences that exemplify the strong but not the weak principle. By their nature, then, the relevant claims will be rather complicated in their *non-normative*

¹¹ This point is from Fine (2005).

elements, which makes it more difficult to say whether they are analytic. Indeed, there may be no sharp line between what is analytic and what is just very unarguably true. That being said, I would agree that

- (5) If Ted is good, then it is impossible to be just like Ted in every descriptive respect and not be good.

is analytic.¹² But cognitive universalism does not in itself predict this. The antecedent entails that there is some kind T, such that Ted is a token of T, and T is good_{KIN}. The consequent, in effect, says that no one, *in any possible world*, is T but not good_{PAR}. But from T's being good_{KIN} in the actual world, it does not obviously follow that T is good_{KIN} in all worlds. And if there is a world where T is not good_{KIN}, then CU allows people there to be T but not good_{PAR}. Before I explain why this does not happen, I will set the stage by discussing a different example, from mathematics.

Syntactically, "47 is a prime number" is in the present tense. So one might think its truth-condition is that 47 is prime *at the time of utterance*. After all, that is the semantical contribution the present tense usually makes. But it seems to me that someone who says

- (23) 47 is prime *today*, but on Thursday it will not be.

is showing clear signs of conceptual confusion. Assuming they understand the rest of the sentence, they apparently do not understand "prime" adequately. In other words, it is plausible that

- (24) If a number is prime, then it is always prime.

is analytic. The explanation is that "prime" is a *tenseless* predicate; it may accept tense syntactically speaking ("was prime", "will be prime"), but it does not really have tense semantically speaking. Exactly how to cash this out depends on your general semantical framework; perhaps the proposition expressed by "47 is prime" does not have a time parameter, or time argument. In any case, the result is this: if you fully understand "47 is prime", then you know that you are not supposed to evaluate it with respect to any particular time. The result is that (23) is analytically false and (24) analytically true; but unobviously so, due to the subtle explanation of their analytic truth/falsehood. That seems about right.

Suppose, next, that someone says

- (25) 47 is prime, but if Gore had won the Florida recount, 47 would not have been prime.

¹² Like before, read "descriptively alike" to include both properties and relations.

That assertion would also seem to signal conceptual confusion. Given a yes/no choice, again, I would say that

(26) If a number is prime, then it is necessarily prime.

is analytic. The explanation is that “prime” is what Fine (2005) calls an “unworldly” predicate. This is the modal analogue of tenselessness. Again, exactly how to cash it out depends on one’s general semantical framework; perhaps the proposition expressed by “47 is prime” does not have a world parameter or world index. In any case, the result is this: if you fully understand “47 is prime”, then you know you are not supposed to evaluate it with respect to any particular world. This makes (25) analytically false and (26) analytically true, but in an even more subtle and unobvious way than before; which again seems like a welcome prediction.

The explanation, I propose, of the analyticity of strong supervenience is cognitive universalism plus the fact that the basic, kind-applying senses of normative predicates are tenseless and unworldly. The analyticity of (5) is then explained as follows. From the antecedent, it follows by UC that Ted is a token of some kind, let us call it “T”, such that T is good_{KIN}. Since “good_{KIN}” is tenseless and unworldly, T is good_{KIN} in every world. But then, by UC, any possible thing that is T is also good_{PAR}, so the consequent is established.

The tenselessness and unworldliness of the kind-applying senses is easy to miss, because the words we use to express them also express particular-applying senses, which are both tensed and worldly. Or more accurately, the particular-applying senses are “bastards”, having a worldly part (the descriptive part) and an unworldly part (the normative) in their definition.

Another reason one might think normative claims about kinds are sensitive to time or world is that we often make such claims in a “parochial” way, leaving some background conditions implicit. For example, we think

(27) Slavery is wrong.

is true and obviously so. But it seems *possible* for slavery to be right; suppose we come across an alien, intelligent species for which freedom is torture and servitude bliss, and which does not even have the potential to be happy in any other way. It seems it would be OK, perhaps a duty, to enslave them.

It is important to notice, however, how we *react* to this kind of example. My reaction is to *clarify* what I mean by (27).¹³ When I now assert (27), I do not intend “slavery” to cover this kind of case; what I intend to say is

¹³ As mentioned in section 10.2 and explained in n. 5, my account can accommodate exception-allowing, generic readings for claims like (27). I cannot feel any such reading of (27), but if you can, restrict what I say in this section to the kind-referring reading.

that *enslavement of humans and other creatures that are capable of autonomy without agony*, is wrong.

There is no point in lengthy clarifications of precisely which act kind we are talking about when our interlocutor understands it perfectly well anyway. There may be bizarre but metaphysically possible scenarios where slavery, rape, genocide, etc., would be OK; but it would only waste time and strain our interlocutor's patience to enumerate outlandish exceptions that are irrelevant to the moral problems we are dealing with.

The crucial point, again, is this: if such a faraway scenario *is* brought to attention and made relevant, we treat it as calling for *clarification* of the claim we made, for example that slavery is wrong. We do not stick to the original claim but reserve it for the actual world; we do not conclude that slavery can be right but is actually wrong.

A careful look at how we react to such examples, then, indicates that the kind-applying normative concepts are indeed unworldly. Which, together with cognitive universalism, explains the analyticity of strong supervenience. I suggest in conclusion that these explanatory advantages, and the structure of moral epistemology, make cognitive universalism an attractive account of normative language.

10.6 THE METAPHYSICAL CHALLENGE

There is another supervenience challenge facing non-natural realism. We now have an explanation of why it is analytic that normative properties supervene on descriptive properties. But, ironically, non-natural realists also have to explain why it is *true* that normative properties so supervene. That might seem confused: analyticity entails truth, so why is there still a problem?

Because what is analytic is that *the properties picked out by normative predicates* supervene on descriptive properties. But it does not follow that *non-natural normative properties* so supervene. That only follows if we assume that normative predicates pick out non-natural normative properties. But if we can give no independent assurance that non-natural normative properties are well-behaved in respect of supervenience, and in fact have some reason to doubt it, then we have no right to this assumption.

Cognitive universalism, in other words, is a theory about English; it explains why the normative part of our language behaves a certain way. Now we need to explain why *the stuff in our theory* behaves in the corresponding way.

Why doubt it? The concern is that non-natural realists claim that normative properties are *fundamental* and non-descriptive. They cannot be reduced in terms of our desires, beliefs, attitudes, or anything else; in fact they are properties of a different *kind* from all others. But on the other hand, given supervenience, they are each going to have a cointensional (necessarily coextensional) 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.¹⁴

This result collides with *Hume's Dictum*: the principle that there be no "necessary connections between distinct existences". It is not obvious exactly what that means (cf. Wilson 2010), but in the present context it boils down to a ban on distinct, cointensional 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, then surely it is possible for *something* to be F without being G or vice versa?

Hume's Dictum is especially plausible for fundamental properties. If F is a fundamental property, then a thing *x*'s being F does not consist in something else being the case with *x*; *x*'s being F is just a basic fact. And likewise for *x*'s being G, if G is also fundamental. But if *x*'s being F is just a basic fact, then there should be another world that is otherwise like the given one, except that *x* is *not* F.¹⁵ And if *x*'s being G is another basic fact, then it would be strange if it were impossible to remove *x*'s F-ness without removing its G-ness. Thus, even looking just at a *single* thing *x*, it seems like it should be possible for fundamental properties to come apart.

The worry, then, is that we have no good explanation of why non-natural 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 non-natural realism is to be plausible, it needs to explain why non-natural reasons or values 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.

¹⁴ See Jackson (1998: ch. 5) for a rigorous statement of this point.

¹⁵ It may be that *x* would need some other property from F's contrast class. For example, if F is a shape, then *x* would take some other shape H. But if G is some other fundamental property (by hypothesis not a shape), then it would be strange if this change would have to result in the removal of *x*'s G-ness.

10.7 MORAL PLATONISM

Many seem to think the non-natural, irreducibly normative properties that Moore, Parfit, Nagel, etc., posit take particular things as their bearers. For example, that Moore posited a non-natural 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 see no reason why the view should take that form. The main intellectual motivation behind non-natural realism has always been a strong commitment to certain pre-theoretical, first-order normative beliefs. And now that cognitive universalism has shown 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, non-natural 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).

The view requires metaphysical commitment not only to non-natural 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. 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 particular-applying property *goodness* will just be the property *having a goodmaking property*. Particular-applying *wrongness* will just be the property *having a wrongmaking property*. And so on. The explanation of supervenience proceeds as before.

I do not mean to suggest that there is a deep metaphysical difference between properties and kinds. There seems to be a shallow difference, marked by the linguistic phenomena above; the kind *murder* is wrong, but the property *murderhood* is not wrong (but wrongmaking). I am not completely sure what to think about this, but it seems to me that the two versions are different ways of spelling out the details of a single underlying view. The underlying view is that the basic normative properties are *second-order*; they take *universals*, rather than particulars, as their bearers. Hence the name “moral platonism”. It is this broad view I want to defend.¹⁶ For ease of exposition, I will revert to the kind-based view, but my discussion applies equally to the property-based view.

Moral platonism 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

¹⁶ The view is suggested by Mackie (1977: 23, 41) and superbly defended by Forrest (1986).

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. Non-natural 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 good places for explanations to end. But it is hard to argue about where explanations should end; so let us just record that the 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 coin-tensional descriptive properties, like goodness and D, as discussed above. But particular-applying goodness is not a fundamental property; it is just the property *being a token of a good kind*. That it 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 between them 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.

There are other examples of necessary connections between first- and second-order universals. Crimson, for example, has the property of being a shade of red. And it has that property in every world, or in every world in which it exists, on an Aristotelian view of universals. 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.

Necessary relations between first- and second-order universals, then, cannot be ruled out as a matter of principle. We have to look at each case and see if a given first- and second-order universal are contingently or necessarily related. In the case of determinables and determinates, it is plausible that the connections are necessary, because they hold in virtue of the nature of the universals themselves. Being a shade of red is part of what crimson is.

The normative case is not like that. It is, for example, not part of what agony is that it is bad. But it is still intuitively plausible that we do not have to check with a given world to see whether agony is bad there. Or whether,

say, *treating another as a mere means* is wrong there. If we keep in mind that we are considering the *basic* normative facts, they are intuitively not contingent.

But why? I suggest the non-natural realist should again look to the semantics for guidance. The kind-applying normative concepts, I argued, are unworldly; the propositions they figure in are not to be evaluated with respect to this or that world. The metaphysical analogue of this is what Fine calls “transcendental” facts; facts that hold, not *in* every world, but *independently* of all the worlds (2005: ch. 10).

Fine makes a distinction, in other words, within the class of what is usually thought of as the necessary truths. The necessary truths proper are those that engage with each world, but in such a way as to come out true every time. “P or not-P” is one example; in some worlds it holds because “P” is true, in other worlds because “not-P” is. Another example is “Nothing is both red and green all over”. Whichever world it is evaluated at, it looks at every object in that world, but never finds a counter-instance.

The transcendental truths, on the other hand, do not even engage the different worlds; there is nothing *in* the worlds that they answer to for their truth. Mathematical claims are good examples: “Seven is a prime number” can only in a degenerate sense be said to be true “in” a given world, for there is nothing about any of the worlds that makes it true.

The distinction between necessary and transcendent truths depends on a distinction between two ways of thinking about possible worlds. On what we can call the *tractarian* conception, a world is a totality of facts. The actual world is everything that is the case; a merely possible world is a way everything could have been. On this conception, there will be no distinction between necessary and transcendental truths; the fact that seven is prime, for example, is a part of every world.

On the second conception, a world is a totality of concrete substances and their attributes. If you specify all the concrete substances there ever are, and all the properties and relations they have throughout their careers, then you have specified a world. Let us call this the *substantive* conception of a world. The actual world, in this sense, is just the physical universe, but substantive worlds need not be physical; some have angels and spooks in them.

The tractarian notion of possible worlds and the corresponding notion of necessity are useful for many purposes, but I suggest that they are not helpful in understanding the content of certain philosophical views. When the mathematical platonist says that mathematical facts are necessary, she is not well understood as saying that they hold, again and again, in each possible world. The mathematical facts hold once and for all, as it were; they belong to an “invariable framework” within which possibilities play out (Fine 2005: 326). To make the point vivid, we can imagine God pondering

which world to create. His alternatives are the substantive possible worlds, not the tractarian ones. Seven being a prime number is not somehow a feature of every alternative, it is not a feature of any of them. The mathematical facts are already there, before he creates anything.

Likewise, when non-natural realists say the basic normative facts are necessary (Parfit 2011: II, 489; Enoch 2011: 146; Scanlon 2014: 39–41), they are not well understood as saying that these facts hold, again and again, in every world. They are part of the invariable framework. God may decide which world to create, but he does not get to say how good it will be if created. The basic normative facts are already there.

The strong supervenience of particular-applying normative properties on descriptive properties is a natural consequence of this view. Whether or not the view violates Hume's Dictum depends on how we understand the Dictum, once the worldly/unworldly and necessary/transcendental distinctions have been made. One could combine moral platonism with a version of the Dictum, restricted to necessary relations between worldly properties (properties that figure in worldly facts), or between fundamental worldly properties.

But a proponent of Hume's Dictum could reasonably say that the spirit of the Dictum requires a ban on both necessary and transcendental connections between distinct properties. In that case the present view contradicts her principle. But someone who is otherwise inclined to accept transcendental facts is not likely to worry that they connect "distinct existences" (what else would they connect?). So Hume's Dictum turns out to be a side issue; the big question is whether to accept transcendental facts in the first place.

Not surprisingly, then, the case for moral platonism opens up into the case for platonist metaphysics in general. If mathematical, logical, and/or modal facts are best understood as transcendental, then moral platonism has impressive allies. I leave it to others to defend the other platonisms; here I will only point out that one can also argue in the other direction. If non-natural realism is the best account of value, and if, as I have argued here, thinking about supervenience takes us from non-natural realism to platonism, then one might simply conclude that there are transcendental facts because there are values.

References

- Blackburn, S. 1971. 'Moral Realism', in J. Casey (ed.), *Morality and Moral Reasoning: Five Essays in Ethics*. London: Methuen, 101–24.

- Blackburn, S. 1985. 'Supervenience Revisited', in I. Hacking (ed.), *Exercises in Analysis: Essays by Students of Casimir Lewy*. Cambridge: Cambridge University Press, 47–68.
- Blackburn, S. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Enoch, D. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Fine, K. 2005. *Modality and Tense*. Oxford: Oxford University Press.
- Forrest, P. 1986. 'Supervenience: The Grand-Property Hypothesis', *Australasian Journal of Philosophy* 66: 1–12.
- Gibbard, A. 1990. *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Gibbard, A. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Hare, R. M. 1952. *The Language of Morals*. Oxford: Oxford University Press.
- Jackson, F. 1998. *From Metaphysics to Ethics*. Oxford: Oxford University Press.
- Koslicki, K. 1999. 'Genericity and Logical Form', *Mind and Language* 14: 441–67.
- Kramer, M. 2009. *Moral Realism as a Moral Doctrine*. Malden, MA: Wiley-Blackwell.
- Leslie, S.-J. 2008. 'Generics: Cognition and Acquisition', *Philosophical Review* 117: 1–47.
- Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Moore, G. E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.
- Scanlon, T. M. 2014. *Being Realistic About Reasons*. Oxford: Oxford University Press.
- Sturgeon, N. 2008. 'Doubts about the Supervenience of the Evaluative', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 4*. Oxford: Oxford University Press, 53–90.
- Wilson, J. 2010. 'What is Hume's Dictum, and Why Believe It?' *Philosophy and Phenomenological Research* 80: 595–637.