

3

FROM PLOTINUS TO RORTY

A History of Philosophy Without Any Gaps¹

Shamik Dasgupta

3.1 Introduction

Wilfrid Sellars famously remarked that philosophy is about ‘how things in the broadest possible sense of the term hang together in the broadest possible sense of the term’ (1962: 35). Perhaps nowhere is this more apt than in contemporary meta-ethics, which revolves around the question of how *values* fit into a world of *fact*.

Admittedly, the question is tricky to state properly (can’t there be facts *about* values?) But details aside, the issue is clear enough. On the one hand there are *material things* – tables and chairs, trees and mountains, stars and galaxies, and indeed everything in the physical universe. If you believe in *immaterial things* like numbers and souls, fine; throw them in too. Call the way *all that* hangs together the ‘descriptive’ realm. On the other hand, we all have opinions about what’s *good* and *bad*, *right* and *wrong*, what *should* be done and the *reasons* for doing it, and so forth; call all *that* the ‘normative’ realm. Then the question is how the normative and descriptive realms hang together.

Many philosophers think that normativity can be located within the descriptive realm. These include utilitarians who think that normative truths ultimately consist in truths about maximizing total happiness; divine-command theorists who think that normative truths consist in truths about the will of a divine intellect; and also expressivists, constructivists, and response-dependence theorists who (in their own distinctive ways) think it all boils down to our normative *attitudes*.² Other philosophers think that normativity cannot be found within the descriptive realm.

These philosophers fall into two camps. One camp infers that normativity is therefore an *extra* component of reality over and above the realm of descriptive fact (they are sometimes called ‘robust realists’). The other camp insists that reality is *purely* descriptive, and so is forced to conclude that normativity isn’t real. Normative talk is *just talk*, on their view – at best, a useful fiction that helps us get along (they are known as ‘error theorists’ or ‘fictionalists’).

What all these philosophers share is the presumption that the descriptive realm cannot itself be located within the normative realm. Think of all the matter spread across space and time. It’s a good question what it ultimately consists in: perhaps it’s just atoms spinning in a void, or the undulations of a quantum wave-function, or perhaps it’s all just ideas in the mind of God. But whatever the final story, the shared assumption is that the descriptive realm is *fundamentally descriptive*: it’s descriptive all the way down. On this picture, the realm of descriptive fact is assumed as given and the question is whether (and how) normativity fits into it.

But an alternative picture is that reality is *fundamentally normative*. The view is that the entire descriptive realm, including all the matter distributed across the physical universe, ultimately consists in pure normativity. This might seem too crazy to take seriously, perhaps even unintelligible. That would explain why this view is never even *articulated* in contemporary meta-ethics, as far as I know, let alone defended!³ But perhaps it has precedence in the ancient writings of Neoplatonists such as Plotinus, who proposed that reality emanates from a single principle of ‘the Good’, with matter appearing at the end of a series of emanations. So I’ll call this alternative picture *normative Neoplatonism* to mark the structural similarity with Plotinus.⁴

Crazy though this view might sound, I believe a surprisingly cogent argument can be made for it. Quite simply, the argument is that *the fate of objectivity depends on it*. Without normative Neoplatonism, I’ll argue, objectivity in both thought and action is a complete non-starter: the idea that some beliefs are objectively *true*, and that some actions are objectively *good*, collapses in a heap, and one is left with the philosophical outlook of Richard Rorty on which the true and the good are not ‘out there’ to be discovered but are made-up by us.

I should say that I am happy to scrap objectivity and to embrace the Rortian vision, so this argument does not convince *me* of normative Neoplatonism. But if you like objectivity, the argument shows that you must accept normative Neoplatonism. Either way, what the argument shows is that you can be a normative Neoplatonist, or you can be a Rortian, *but those are your only options*. There is no in-between.⁵

3.2 Objectivity in Action

We must start by identifying the sense of objectivity at stake. To this end, imagine coming across a hitherto isolated community of humans that live a very different form of life from us. They engage in activities that we consider reprehensible, including ritual human sacrifice. And they believe a strange origin story of being transported from the spirit world on the back of a flying pig. We think they are deeply misguided, both in their *actions* (of human sacrifice) and *beliefs* (in flying pigs). But they presumably think the same about us, and there's a kind of 'cultural relativism' that says the following:

Look, they have their form of life – their way of navigating social interactions, the narratives they tell about the world – and we have ours. We call our narratives “science” and theirs “myth”, our social codes “progressive” and theirs “primitive”, but that's just arrogance. In truth, neither form of life is inherently more valid than the other; they're just different.

The objectivity at issue is the *antithesis* of this and holds that some forms of life are 'objectively' more valid than others – more valid, that is, not by the lights of one or other community, but by the lights of reality itself (so to speak). But what does this really mean?

To answer this, it will help to clarify what relativism is in the first place. To this end, focus just on *action* for now, and, in particular, on the evaluation of actions as *right* or *wrong*. According to a simple relativist model, there are no facts about which actions are right or wrong *simpliciter*, just facts about which are right or wrong *relative to a given ethical code*, where all codes are on a par. Suppose we accept a utilitarian code that recommends actions that maximize total pleasure. And suppose they accept a divine-command code that recommends actions that promote the will of their God, as expressed in some revered text – which, let's suppose, calls for the human sacrifice of atheists. Then killing atheists is right relative to *their* ethical code; feeding atheists is right relative to *our* ethical code (let's suppose); and according to relativism that is all there is to say.

If that's relativism, what must be the opposing *objectivist* view? At a minimum, it must hold that one community can be legitimately criticized as acting *wrongly*. But that is not enough, since a sophisticated relativist can allow for this by noting that we can evaluate other cultures from the perspective of *our* ethical code, not *theirs*. More fully, the idea is that a statement evaluating an action as right (or wrong) is true if the action is recommended (or prohibited) by the ethical code of *the person making the*

statement, not the person doing the action. Thus, when we evaluate the other community and say

(1) Killing atheists is wrong and feeding them is right,

that is *true* (in our mouths), according to sophisticated relativism, because *our* utilitarian code prohibits killing atheists and recommends feeding them. Yet this is not enough to secure the idea that our way of life is ‘objectively’ more valid than theirs, because of course the same goes in reverse: when the other community evaluates *us* and says

(2) Feeding atheists is wrong and killing them is right,

that is *also true* (in their mouths) because *their* ethical code prohibits feeding atheists and recommends killing them instead. Thus, on this sophisticated relativism, we can legitimately criticize them as doing wrong; the trouble is that they can *with equal legitimacy* criticize *us* as doing wrong in return!

The objectivist must somehow break this ethical symmetry, but how? The obvious idea is to drop the relativization to ethical codes altogether and insist that an action can be right (or wrong) *simpliciter*. The objectivist then owes us an account of what makes an action right *simpliciter*. It doesn’t matter for our purposes what the account is, but for concreteness let’s suppose that an action is right *simpliciter* if and only if it maximizes total pleasure. Note that there is no relativization to ethical codes here; rather, the proposal is that utilitarianism is the *true and universal* account of right action. It follows that (1) is true and (2) is false *simpliciter*, whether spoken in our mouths or theirs. We act rightly and they act wrongly, *period*.

Is this enough to secure the idea that our way of life is ‘objectively’ more valid than theirs? Unfortunately not. To see why, note that all the objectivist has said is that the concept ‘right’ applies to all and only those actions with the property of maximizing total pleasure. In effect, this property is the standard by which we evaluate actions when deciding what to do. But of course, there is also the property of promoting the will of the other community’s God, and so there is nothing to stop that community from introducing a *different* concept – call it ‘right*’ – that applies to all and only those actions with this *other* property, and then using it as the standard by which *they* evaluate actions when deciding what to do. Then, while

(1) Killing atheists is wrong and feeding them is right

is true *simpliciter*, the statement

(2*) Feeding atheists is wrong* and killing them is right*

is *also true simpliciter*! That is, even if pleasure-maximization is the true and universal account of *right* action, promoting their God’s will

is the true and universal account of *right** action. And now we have symmetry regained, for nothing the objectivist has said so far distinguishes the property *we* organize our lives around and call 'right', over the property *they* organize their lives around and call 'right*'. Thus, we can criticize them by saying (1), and that would be true; the trouble is, they can *with equal legitimacy* criticize us by saying (2*) – and that would *also be true!*

I've focused on 'right' and 'wrong', but the same applies to other concepts we use to evaluate our practical lives such as 'good', 'bad', 'should', 'reason', 'obligation', 'virtue', and so on. In order for our way of living to be 'objectively' more valid than the other community's, it's not enough that these concepts express non-relativized standards involving (say) pleasure-maximization, such that one *should* feed atheists, or that there is more *reason* to feed them than kill them, etc. All that may be true – and true *simpliciter*. The trouble is, it will be *equally true* – and true *simpliciter* – that one *should** kill atheists, that there is more *reason** to kill them than feed them, and so on, where these *-ed concepts express non-relativized standards involving the will of their God. Hence the symmetry, for nothing so far has been said to distinguish the (non-relativized) standards *we* express with our concepts and organize our lives around, over the (non-relativized) standards that *they* express with their *-ed concepts and organize *their* lives around.

You might try saying that one *should* live by our standards, not theirs, or that there is *reason* to prefer our concepts over theirs; but that would be futile. For it will be equally true that one *should** use their standards, not ours, and that there is *reason** to prefer their concepts over ours. That is, if our way of life is preferred *by the lights of our concepts* 'should' and 'reason', their way of life will equally be preferred *by the lights of their concepts* 'should*' and 'reason*'. What we haven't yet captured is the objectivist idea that one way of life is preferred *by the lights of reality itself*. Both ways of living remain on a par.

What objectivists need is a symmetry-breaker – something that distinguishes, say, *our* concepts and standards over *theirs* – and (as we just saw) they won't find one in our received conceptual repertoire. Objectivists must therefore posit something new – a further fact that distinguishes one standard as *normatively significant* over the rest. Note that 'normative significance' should not be understood in its colloquial sense, for the above reason; it is, rather, a new theoretical term introduced with the stipulation that it expresses the further fact that breaks the symmetry (if there is such a fact). If a standard S is normatively significant, then *by definition* S determines the way of life preferred not just by the lights of this or that community, but by the lights of reality itself. It determines what's Really-Right, as I'll put it, where this may in principle differ from

what's 'right' in our colloquial sense; it determines the Real-Reasons and Real-Obligations, and so on. It follows that normative significance is not *-able. If our standard of pleasure-maximization is normatively significant, the other community could of course introduce the *words* 'normative* significance*' to describe their standard of promoting God's will. But that would not *genuinely* restore symmetry, for the property of normative significance (if there is such a thing) is by definition a symmetry-breaker. 'Normative* significance*' would be a mere symmetry in language, as it were, not in ethical reality.

Crucially, normative significance must be *primitively normative* – that is, not reducible to merely descriptive properties. For suppose otherwise: suppose the standard of pleasure-maximization, say, was normatively significant in virtue of having some descriptive property P such as being the standard *we* organize our lives around, or of having such-and-such descriptive consequences. Well, there will of course be a descriptive property G that distinguishes the standard of promoting their God's will, such as being the standard that *they* organize their lives around, or of having so-and-so consequences. So, for pleasure-maximization to be normatively significant over their God's will, P must *already* have normative significance over G. If this is an irreducible fact about P, we're done. Alternatively, if this fact about P holds in virtue of some further descriptive property P* that P has, there will be a corresponding descriptive property G* that G also has ... and so on. At some point *something* must have normative significance on its own (that is, not in virtue of anything else), or else the chain of normative significance goes on forever. Either way, we have primitive normativity – a normative property not reducible to merely descriptive properties.

The question was what it could *mean* for some actions to be 'objectively' more valid than others. And the answer is that there must be primitive normativity – an irreducible property of *normative significance* that distinguishes one standard of evaluating actions over the others. Without that, we're stuck with the cultural relativism we started with on which all ways of living are 'equally valid': they will each satisfy *some* standard or other, and there will be no further fact as to which standard is 'correct' in a universal, objective sense. As Richard Rorty said, we will have to 'give up on the idea that there are unconditional, transcultural moral obligations, obligations rooted in an unchanging, ahistorical human nature' (1999: xvi). To avoid this Rortian conclusion and secure moral objectivity, we must embrace primitive normativity.

3.3 Objectivity in Belief

So much for *action*; what about *belief*? The other community believes an origin myth that involves flying pigs, yet the cultural relativist said that their beliefs are ‘no less inherently valid’ than ours. Their myths are just as correct *for them* as science is *for us*. I said that the objectivity at issue is the *antithesis* of this and holds that some beliefs are ‘objectively’ more correct than others – correct not just by the lights of one or other community, but by the lights of reality itself. But what does this mean?

Well, what could the *relativist* mean by saying that the other community’s beliefs are ‘just as valid’ as ours? It can’t just be that the other community has less experience with pigs, so that their belief in flying pigs is rational *given their limited evidential basis*. For it’s a truism that what one should rationally believe depends upon one’s evidence: relativists and objectivists agree on that. Thus, let’s imagine that the other community has the same evidential basis E as we do: they’ve never seen flying pigs, they’ve seen what happens when pigs fall from great heights, etc. This evidential basis E obviously justifies the belief that pigs *can’t* fly – or, at least, it does so by the lights of *our* standards of rationality S. So, the other community must use *different* standards of rationality, S*, relative to which E justifies the belief that pigs *can* fly. At a minimum, then, relativism must be the view that S and S* are on a par: neither is ‘the correct’ standard of rational belief. Objectivists will presumably insist that the other community’s belief is *false* and ours *true*, and that our standards S are distinguished over S* in virtue of being better at tracking truth. So, relativists must add that there is no such thing as truth and falsity *simpliciter*, just truth and falsity *relative to an epistemic standard*: the other community’s beliefs are true *relative to their standards*, our beliefs are true *relative to our standards*, and there is no further fact of the matter as to whose beliefs are ‘really’ true.

If that’s relativism, what must be the opposing ‘objectivist’ view? At a minimum, it must hold that we can legitimately criticize the other community’s beliefs as *false*. But that is not enough, for (as before) a sophisticated relativist may say that we can evaluate their beliefs from the perspective of *our* standards, not theirs. If so, we can indeed say that their belief that pigs fly is false (relative to our standards); the trouble is that they can *with equal legitimacy* say that *our* belief that pigs can’t fly is *also* false (relative to *their* standards)!

Objectivists must somehow break the symmetry and insist that beliefs can be true or false *simpliciter*. The rough idea is that a true belief is one that reflects reality, or ‘corresponds to the facts’ – no relativization to rational standards involved. This idea can be developed in numerous ways,

but a natural approach starts from the idea that there is a relation of *representation*, or ‘reference’, between words and reality. Thus, in

(3) Birds fly,

the word ‘birds’ refers to (or represents) a certain kind of animal, namely *birds*; ‘fly’ refers to (or represents) a certain kind of activity, namely *flight*; and (3) is true because the animal referred to by ‘birds’ engages in the activity referred to by ‘fly’. Note that there is no relativization to epistemic standards here: (3) is true *simpliciter*, thanks to representing (or ‘corresponding to’) reality. And what goes for *sentence* (3) goes for the *belief* that one would verbalize with it: our belief that birds fly is also true *simpliciter*, for the same reason.⁶ Generalizing beyond (3), the idea is to define a non-relativized property of truth in terms of reference along the following lines:

A sentence of the form ‘*a* is *F*’ is *true* if and only if the referent of ‘*a*’ has the feature referred to by ‘*F*’.⁷

This implies, for example, that

(4) Pigs fly

is true if and only if the referent of ‘pigs’ (namely, *pigs*) engages in the activity referred to by ‘fly’ (namely, *flight*). Which they don’t; hence (4) is *false* – again, not false relative to an epistemic standard, but false *simpliciter*.

Is this enough to secure the idea that our belief in (3) is ‘objectively’ more valid than the other community’s belief in (4)? Unfortunately not. To see why, note that all that we’ve been given is a definition of truth in terms of *reference*, a relation that holds between ‘birds’ and birds. But the word ‘bird’ stands in *all sorts* of other relations to other worldly items. After all, what is a relation? Abstractly, a relation can be represented as a set of pairs of objects. For example, the relation of *loving* is a set containing

<Romeo, Juliet>,

<Me, my wife>,

...

and so on, for every pair of things such that the first loves the second. Likewise, the word-world relation of *reference* is a set containing

<‘birds’, birds>,

<‘pigs’, pigs>,

<‘flies’, flight>,

...

and so on. But sets of pairs are plentiful; there is also a set consisting of the following pairs:

<'birds', pigs>
 <'pigs', birds>
 <'flies', flight>
 ...

This set of pairs corresponds to a *different* word-world relation: call it *reference**. Thus, our word 'pigs' *refers* to pigs and *refers** to birds *at the very same time!* If this sounds bizarre, consider that I am both *the son of Carol* and *the father of Dylan* at the very same time. There is nothing strange about this; it's just the mundane fact that I stand in many different relations to many different things at once. My point is that the same goes for relations between words and world.

Recall that on the current view, truth is just a property of sentences defined in terms of the word-world relation of *reference*. Since there is also the word-world relation of *reference**, there will be a corresponding property of *truth** defined analogously in terms of *reference** as follows:

A sentence of the form 'a is F' is *true** if and only if the referent* of 'a' has the feature referred* to by 'F'.

It follows that while (3) is true and (4) is false (*simpliciter*), it is *also* the case that (4) is true* and (3) is false* (*simpliciter*)! And now symmetry is regained, for nothing the objectivist has said so far distinguishes truth over truth*. Thus, we can boast that our beliefs are *true* and criticize theirs as *false*; the trouble is, they can *with equal legitimacy* boast that their beliefs are *true** and criticize ours as *false**! To be sure, we may not *care* much about the property of truth*: we presumably aim to believe things that are *true*, not true*. But, equally, they presumably don't care about truth; they aim (let's suppose) to believe things that are *true**, not true! For all that's been said, both properties are on a par.⁸

Objectivists might now try saying that one *should* aim for truth (not truth*), or that true beliefs are the *right* way to represent the world. But that would be futile, for it just uses *our received normative concepts* to evaluate truth and truth*, and the same dialectic will play out again: the other community can with equal legitimacy say that one *should** aim for truth* (not truth), and that truth* is the *right** way to represent the world, where their *-ed normative concepts stand to truth* just as our normative concepts stand to truth. The upshot is that while our beliefs are preferred *by the lights of our concepts* of 'truth' and 'should' and 'right', their beliefs

are preferred *by the lights of their concepts* of ‘truth*’ and ‘should*’ and ‘right*’. What we have not yet captured is the objectivist idea that our beliefs are preferred *by the lights of reality itself*.

What objectivists need is a symmetry-breaker, and (as before) they won’t find one in our received conceptual repertoire. They must therefore posit something new, a further fact that distinguishes, say, reference and truth as *normatively significant* over reference* and truth* – a fact in virtue of which *true* beliefs are the ‘Really-Right’ way to represent the world, or ‘Really-True’, as I’ll put it. Again, ‘normative significance’ is not used in its ordinary English sense here; it is a new, theoretical term introduced with the stipulation that it expresses the further fact that breaks the symmetry (if there is such a fact). And for the same reason as before, this property of being normatively significant must be *primitively normative*, not reducible to the descriptive properties of whatever standard possesses it.

The question was what it could *mean* for some beliefs to be ‘objectively valid’ – correct not just by the lights of one or other community, but by the lights of reality itself. And the answer is that there must be a primitive, irreducible property of *normative significance* that distinguishes one standard of evaluating beliefs over the others. Without that, we’re stuck with the cultural relativism we started with, on which our ‘science’ is just another narrative on a par with their ‘mythology’. As Richard Rorty put it, ‘no organism, human or non-human, is ever more or less in touch with reality than any other organism’; hence ‘there is no point in asking whether a belief represents reality, either mental reality or physical reality, accurately’ (1999: xxiii–iv). To avoid this Rortian picture and to secure the notion of *objectively* correct representations of reality, we must embrace the property of normative significance just described.⁹

3.4 The Interaction Problem

I’ve argued so far that *objectivity requires normative significance*. This goes for both action and belief. For an action to be objectively more valid than others, it’s not enough that it satisfies a standard we call ‘right’, such as pleasure-maximization; it must also be that that standard is normatively significant. And for a belief to be objectively more correct than others, it’s not enough that it is verified by just *any old* word-world correspondence relation; it must also be that that correspondence relation is normatively significant.

What I’ll now argue is that normative significance is highly problematic ... *unless* you’re a normative Neoplatonist! If that’s right, we have our argument for normative Neoplatonism: anyone who likes objectivity should endorse it.

The key claim here is conditional – namely, that normative significance is problematic *if* normative Neoplatonism is false. To see why, let's assume the antecedent. Assume, that is, that the descriptive realm is *fundamentally descriptive*: descriptive facts about, say, the distribution of matter throughout spacetime do *not* hold in virtue of purely normative matters. Likewise for descriptive facts about how people happen to act and about which standards they satisfy, such as that we organize our lives around the standard of pleasure-maximization and call actions that satisfy it 'right', and that the other community uses the standard of promoting their God's will and calls actions that satisfy *it* 'right*'. This is all in the realm of descriptive fact, which (we're assuming) is fundamentally descriptive and does *not* consist in a purely normative basis.

Why, then, would normative significance be problematic? To explain, I'll focus on the case of action (the argument generalizes straightforwardly to the case of belief). If a standard S is normatively significant with respect to action, two things follow. (a) By definition, S determines which actions are 'objectively' correct – correct not just by the lights of the parochial standards of this or that community, but by the lights of reality itself. It determines what's Really-Right, the Real-Obligations and Real-Reasons, and so forth. (b) At the same time, we saw that there is nothing in virtue of which S plays this role: it must be an inexplicable, brute fact that S, not some other descriptive standard, has this property of normative significance. Now, look at how bizarre this is. By (a), S is a *remarkable* bit of reality. It fixes reality's preferred way of living, regardless of the various parochial standards that we've come to adopt. Surely, you would think, there must be some explanation of why S is so special, something about it that makes it fit to play this extraordinary role. But by (b), there isn't. Nothing explains why reality's preferred way of acting springs from *it* rather than anything else; it's simply a brute fact about reality that they do. To see how bizarre this is, suppose someone suggests that S is my daughter's will. Their view is that the Really-Right thing to do is whatever promotes *her* will, not God's – that the Real-Reasons and Real-Obligations and so forth all align around *her*. It follows that my daughter is *really* special: it's her will, not yours or mine, that fixes reality's preferred way of living! But if we ask what *makes* her so special, there's no answer: there's nothing about her that makes her more fit than you or me to play this role. And that seems ludicrous. What's ludicrous here is here not *just* that the preferred way of living is fixed by her (though that certainly is ludicrous). It's that the preferred way of living could be so *arbitrary*. It's *reality's preferred way of living*, after all – something far too important for there to be no rhyme or reason as to why it is what it is.

So far, this is just an ‘intuition’ that something is odd, and intuitions can be fallible. But in this case, I believe the oddity is indicative of a deeper problem – namely, that the view in question is (in effect) a *dualist* view of the descriptive and normative realms. Think of it on the model of Descartes’s mind-body dualism. On Descartes’s view, one’s body is one thing and one’s mind another; the latter is an *immaterial* substance that does not exist in space. Princess Elisabeth famously raised the ‘interaction’ problem: if one’s mind is immaterial, how can it have an effect on one’s body? How can *thinking* cause *action*? I want to raise a similar problem. The view that I’m criticizing is a *normative-descriptive* dualism on which the descriptive realm is one thing and the normative realm another, and the problem is how these two realms come to be linked up in the way they are.

To see the point, let’s examine the interaction problem for mind-body dualism in more detail. In fact, there are two related problems. One is *mechanistic*, of how an immaterial substance could come into *causal contact* with matter. But a different problem is *correlational*, of how mental events come to be correlated with material events in the way they do. Princess Elisabeth focused on the first problem:

It appears that all determination of movement is produced by the pushing of the thing being moved, by the manner in which it is pushed by that which moves it, or else by the qualification and figure of the surface of the latter. Contact is required for the first two conditions, and extension for the third.¹⁰

Her thought is that causation is a mechanistic process that requires pushes and pulls between things in space; hence not something in which an immaterial mind can participate. But this view of causation is not mandatory. On so-called ‘Humean’ accounts, causation is just regularity: for one event to cause another is for events like the first to be followed by events like the second. Here there’s no need for a mechanical process connecting the events; it’s enough that the events unfold in regular patterns. So, as Bennett (2007) emphasized, dualists can say that certain events in the mental realm – say, intentions to raise one’s arm – tend to be followed by certain events in the material realm – say, raisings of one’s arm – and *that’s* all that mind-body causation consists in. The mechanistic problem is therefore not so pressing – at least, not for Humeans.

Still, that leaves the *correlation* problem. To see what this is, suppose my arm raises whenever my immaterial mind is in a certain state P. Then the question is this: why does P correlate with *my arm raising*, not (say) my leg kicking? Why is the raising of my arm always preceded by P, not some other mental state Q? The correlation problem is that this alignment is

brute – an arbitrary connection between two separate domains. The problem is easy to miss because we would normally describe P as an ‘intention to raise one’s arm’, and what is mysterious about the fact that *intending* to raise one’s arm correlates with *raising* it? But this just buries the problem in language. To keep track of things, we must ask why we describe P like that in the first place. Presumably, it has something to do with its *effects*: a mental state counts as an ‘intention to raise one’s arm’ only insofar as it is typically followed by an *arm raising*; if it were typically followed by a *leg kicking*, we would instead describe it as an ‘intention to kick one’s leg’ instead. So, there’s obviously no mystery why a mental state correctly described as ‘an intention to raise one’s arm’ is typically followed by an arm raising; that’s guaranteed by how we describe the mental states. Instead, the mystery concerns the mental states lying behind these descriptions. There are mental states of the kind P and mental states of the kind Q; the question is why the *former* and not the latter are typically followed by a hand-raising. Mind-body dualists appear to have no answer to this. There are *all sorts* of ways in which states of the immaterial realm *could in principle* line up with states of the material realm, and the particular way in which they *do* align appears to be arbitrary – a brute and inexplicable line drawn between two entirely independent domains. That’s the correlation problem.¹¹

The correlation problem would go away, notice, if there was a *non-Humean* solution to the mechanistic problem. For presumably a mechanistic account of how mental states come into *causal contact* with one’s body would explain why the mental states *correlate* with bodily movements in the way that they do: causation implies correlation, after all. But without a mechanistic account, the correlation problem looks serious. For the dualist is telling us that there are two independent realms, the material and the immaterial; and that there are systematic correlations between the items in each realm – but that it is *completely arbitrary* that the two realms align in the particular way they happen to. This kind of brute connection between two distinct realms is metaphysically embarrassing. If nothing else, it means that immaterial minds do nothing to *explain* bodily behavior, for how can an immaterial state P explain the material raising of my arm if the correlation between the two is brute and arbitrary in the first place? Immaterial minds therefore appear to be explanatorily otiose (at least with respect to bodily movement).

Go back now to the issue of normative significance. We’re assuming that the descriptive realm is fundamentally descriptive. Suppose, in addition, that one standard for action is normatively significant over others. Then the result is a dualist view on which the descriptive realm is one thing and the normative realm another. The descriptive realm includes all the

matter spread out across space and time, as well as facts about what people do and about which standards their actions satisfy (as mentioned above). By assumption, all of this holds independently of anything normative. The normative realm, by contrast, includes facts involving Real-Obligations, Real-Reasons, and so forth. Exactly what these are doesn't matter for our purposes, but contenders might include these:

An action is Really-Right if it maximizes what's Really-Good,
 An agent Really-Should do whatever they have most Real-Reason to do,

and so on (just think of familiar normative principles and prepend 'Really-'). Importantly, these leave open *what* one Really-Should do: for example, maximize pleasure or promote God's will. For that wouldn't be a purely normative fact: it would be a 'mixed' fact linking the descriptive and normative realms.¹² Clearly, there are *all sorts* of ways in which the two realms *could in principle* line up. The particular way in which they *do* align depends, of course, on which standard in the descriptive realm is normatively significant. But since normative significance is primitive, what this amounts to is *a brute and inexplicable connection between two entirely independent domains*. This is the correlation problem once again! It is widely seen as an embarrassment for mind-body dualism, and I claim it is no less embarrassing for the normative-descriptive dualism under discussion.¹³

Why is it so embarrassing? In the case of mind-body dualism, one problem was that the mind is left *explanatorily* otiose: an immaterial state P cannot explain why I raised my arm if the correlation between the two states is brute and arbitrary in the first place. But with normative-descriptive dualism the problem is different, for it's unclear whether normativity is *supposed* to explain anything descriptive in the first place. What, then, is normativity supposed to do? It's supposed to make *demands* of us, for one thing. To recognize that an action is the *right* thing to do (or what *should* be done, or what there's most *reason* to do, etc.) is to recognize a sense in which it *must* be done. And that's just parochial normativity: to recognize that an action is *Really-Right* is to recognize an even greater demand, one that springs *from reality itself* independently of what our parochial concept of 'right' demands of us. It is this that the correlation problem exposes as an embarrassment, for it shows that it is *completely arbitrary* which actions are Really-Right: it depends on which standard is normatively significant, yet it's a brute and arbitrary fact which standard that is. And it is hard to see how something so arbitrary, something completely lacking in rhyme or reason, could place *demands* on us in any serious sense. If reality demands that you leave your family and die for your country, and you ask

‘Why?’, the answer cannot be ‘There is no answer’: it cannot be a brute fact that you must sacrifice everything you care about.¹⁴

So, the difference between the two correlation problems is this: in the case of mind-body dualism, it leaves the mental realm unable to *explain* anything about us, while in the case of normative-descriptive dualism, it leaves the normative realm unable to *demand* anything of us. The mind-body correlation problem would go away, we saw, with a ‘mechanistic’ account of how the two realms come into causal contact. But with normative-descriptive dualism there is no such hope, for the problem is not *causal* in the first place! In this respect, the dualism between normativity and description is at least as problematic as the dualism between mind and body.

3.5 Normative Neoplatonism

How does one avoid these interaction problems? By rejecting the dualism that gives rise to them, of course! We must collapse one of the two domains into the other.

With mind and body, one option is to embrace *materialism* and say that mental states ultimately consist in brain states. The correlation between mental states and bodily movements is then a correlation between things within the material realm. Another option is to embrace *idealism* and say that material things ultimately consist of mental ideas, so that the correlation between mental states and bodily movements is really a correlation within the mental realm. There may be further questions about correlations within each realm, of course, but at least there is no problem of how two metaphysically independent domains come to be aligned in the way that they do.

With normativity and description, our options are analogous. One is to say that reality is fundamentally descriptive, so that normative truths ultimately consist in purely descriptive facts of some kind or another. But the problem with this option is that it rules out objectivity, which, as we saw in Sections 3.2 and 3.3, requires *primitive* normativity. The other option is to say that reality is fundamentally normative. *This is normative Neoplatonism*, the view that the entire descriptive realm – including all the matter distributed across the physical universe – ultimately consists in pure normativity. On this view, ordinary descriptive truths like

(5) There is a table in my office

are grounded in purely *normative* truths of some sort. The advantage of this option, of course, is that it makes room for objectivity: the primitive normativity required for objectivity is no problem since all of reality is fundamentally normative anyway.

I promised you an argument for normative Neoplatonism, and we now have all the pieces in place. It goes like this:

- i There are objectively correct things to do and believe. (Premise)
- ii Objectivity requires primitive normativity. (Sections 3.2 and 3.3)

Therefore,

- iii There is primitive normativity. (From (i) and (ii))
- iv Primitive normativity implies normative-descriptive dualism, *unless* the descriptive realm is fundamentally normative anyway, *per* normative Neoplatonism. (Section 3.4)
- v Dualism is false. (Section 3.4)

Therefore,

- vi Normative Neoplatonism is true. (From (iii), (iv), and (v))

Note that (i) is an unargued premise, and I said at the beginning that I'm happy to give it up. I think that reality is fundamentally descriptive, so *for me* the upshot of all this is that objectivity is a will-o'-the-wisp. But I don't mean to downplay its attraction: the idea that there are objectivity correct things to believe and do was a cornerstone of my worldview for the first forty years of my life! I sympathize with those who remain seduced by the idea; my point is just that it commits them to normative Neoplatonism.

Normative Neoplatonism is a wild view, to be sure. It maintains that *normativity* gives rise to the distribution of matter throughout spacetime, and one might well wonder how this is possible – or even whether it's *intelligible*. But to play devil's advocate, note that this is a problem for *all* non-dualist views under discussion. Take the materialist view that mental states consist in brain states. This faces the notorious 'hard problem' of consciousness – the problem of explaining how a material brain could give rise to subjective *feels*, the *what it's like* to be you – and to my mind this problem remains intractable. Likewise, idealists struggle to offer a convincing account of how mental ideas could possibly give rise to, say, a *rock*. If we're honest, we have no idea as to how mind *or* matter could give rise to the other. Materialism and idealism should therefore be seen as black boxes: we can argue for them on the basis of high-level structural virtues, but at the end of the day we don't really know how they work 'inside'. I suggest that we see normative Neoplatonism in the same light.

To play devil's advocate further, perhaps the view is not as inscrutable as it first appears. It's often said that pain is essentially something *bad* – that

it lies in its nature *to be avoided*. The idea is that these normative aspects are not accidental to the state of pain, but instead help to define *what it is*. If so, we have here (the beginnings of) a *definition* of a descriptive state in normative terms – precisely the kind of thing that normative Neoplatonism needs. I've no idea how this might generalize to other descriptive matters, I've merely argued that objectivity is impossible unless it does.¹⁵

Notes

- 1 With apologies to Peter Adamson, whose podcast *The History of Philosophy Without Any Gaps* is an extraordinary philosophical resource from which I first learnt about Plotinus.
- 2 To be clear, some utilitarians will reject this reduction of normativity to descriptive truths about happiness. Still, the utilitarian reduction is a particularly clear model of how a reduction might go. Likewise for divine-command theory.
- 3 Even John Leslie's (1979) 'extreme axiarchism' doesn't count, since it admits as fundamental various mixed truths that involve both normative and descriptive elements; see note 12 below.
- 4 There is some debate as to whether Plotinus' principle deserves a normative title (or any title, for that matter). The worry is that his principle is supposed to be a *unity*; yet to predicate anything of it, such as 'goodness', would result in *two* things – the principle and its goodness. But putting that aside, the principle must be normative in *some* sense because of its role in Plotinus' theory of evil. On that theory, evil consists in the *absence* of emanation – evil begins where emanation from the principle ends – and hence the principle must be normative in a sense.
- 5 In truth, I cannot hope to support a conclusion of this strength in a chapter of this length. What follows is just a sketch of a line of thought that I am developing elsewhere.
- 6 I am slurring over some differences between thought and language, but those complexities won't matter for our purposes.
- 7 This is just a sketch of a base clause for a so-called 'Tarskian' definition of truth.
- 8 The idea that there are alternative world-world relations of reference* and truth* was explored by Stich (1990). And the 'permutation' argument sketched in the text is, of course, related to the model-theoretic argument of Putnam (1977).
- 9 To clarify, the property posited here to secure the objectivity of *belief* need not be the same as the property posited in the previous section to secure the objectivity of *action*. So, it's more accurate to call the former the property of normative significance *with respect to belief*, and the latter the property of normative significance *with respect to action*. But I'm slurring over this detail, for ease of prose.
- 10 Princess Elisabeth in a 1643 letter to Descartes; quoted by Bennett (2007: 318).
- 11 Compare Jaegwon Kim's (2005) pairing problem.
- 12 This is why Leslie's 'extreme axiarchism' (1979) doesn't count as normative Neoplatonism. It takes, as fundamental, mixed facts of the form 'it would be *best* if the world had *such-and-such descriptive feature*'.
- 13 Some believers in primitive normativity argue that it would supervene on descriptive matters. Even so, the point remains that their (supposed) necessary connection would be brute and inexplicable.
- 14 Compare Korsgaard's (1996) 'normative question'.
- 15 Thanks to David Builes and Russell McIntosh for their insightful comments on an earlier draft of this material.

References

- Bennett, Karen (2007). Mental Causation, *Philosophy Compass* 2: 313–37.
- Kim, Jaegwon (2005). *Physicalism, or Something Near Enough*, Princeton, NJ: Princeton University Press.
- Korsgaard, Christine (1996). *The Sources of Normativity*. Cambridge, UK: Cambridge University Press.
- Leslie, John (1979). *Value and Existence*, Totowa, NJ: Rowman & Littlefield.
- Putnam, Hilary (1977). Realism and Reason, *Proceedings and Addresses of the American Philosophical Association* 50: 483–98.
- Rorty, Richard (1999). *Philosophy and Social Hope*, London: Penguin Books.
- Sellars, Wilfrid (1962). Philosophy and the Scientific Image of Man, in *Frontiers of Science and Philosophy*, (ed.) Robert G. Colodny, Pittsburgh: University of Pittsburgh Press: 35–78.
- Stich, Stephen (1990). *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*, Cambridge, MA: The MIT Press.