

## SUBJECT-IVE AND OBJECTIVE

Peter Railton

### I

Morality is different – but how? Moral judgement certainly differs in various ways from ordinary judgement about the garden-variety objects in the world around us. A distinguished line of philosophers have argued that these differences are crystallised in a difference in the very nature and function of moral language. And yet grammar textbooks make no mention of a systematic distinction between moral and non-moral language, even at the level of deep structure. Indeed, moral language does not in its grammatical or logical behaviour betray the least sign of being anything other than ordinary.

Increasingly, philosophers are inclined to take this at face value. Moral judgements, for example, are held to make assertions that can be true or false in our ordinary sense of these terms. Increasingly, too, philosophers are inclined to accept all the platitudes that follow from this. Thus, since ‘truth is correspondence to the facts’, moral truth is correspondence to the facts, too. Ordinary discourse supports us here: ‘The fact is that I’ve promised to go and so I’m obliged to be there’ is not a quirky way of expressing oneself. It even is deemed harmless to countenance talk of ‘moral properties’ (in ordinary speech one would more likely hear ‘moral qualities’) and ‘moral facts’ (‘moral circumstances’). This sort of talk seems to assume nothing about whether such properties or facts are *sui generis*, for example.

Perhaps, then, morality is different but moral discourse as such is not. And perhaps our willingness to talk in terms of moral properties and facts simply registers our readiness to apply the terms ‘true’ and ‘false’ to moral judgements. As such, it might be considered innocent of any significant metaphysical implications. Its real basis would be seen as lying in our discursive practices in morality, our tendencies to deliberate and disagree in distinctive ways – to make arguments, to insist upon consistency, to adduce evidence, and so on. Beyond this, no purported insight into the real composition of the world is claimed.

If moral judgement is therefore factual in a thorough-going but

mundane sense, where might the supposed difference between moral judgements and ordinary worldly judgements lie?

Let us consider a recent suggestion. It is one thing to judge, another to reify the object of judgement. We must always be on guard when analysing language – moral and non-moral alike – against gratuitous reification. Sometimes reification strikes us as quite appropriate. For example, most of us are realists about the external world of mid-sized physical objects, about other minds, and about the past. But many find reification and realism suspect in other domains that nonetheless are typically accepted as factual – mathematics, modality, universals, and so on.

The question of realism vs. anti-realism has therefore seemed to some a promising way to understand how morality might be different even though moral discourse is factual. There is, after all, a lot of ground between recognising that a discourse is factual and adopting a realist attitude toward the substantives and properties occurring in it.<sup>1</sup> First, there is always the possibility that a *non-literal* interpretation of these terms will seem most appropriate: philosophical behaviourists argue, for example, that mental statements are equivalent to claims about actual and possible behaviour. This translation would preserve the cognitive grammar of mentalistic discourse and yield a domain of potential psychological facts, but it would hardly deserve the name ‘mental (or psychological) realism’. Second, there is the further possibility that an *error theory* about the discourse will seem on reflection most appropriate: moral statements when literally interpreted, might be seen as making claims or presuppositions – such as the existence of a teleological ‘natural order’ – that simply never hold true in the actual world.<sup>2</sup> In error theories we encounter a form of anti-realism that owes nothing to a non-factualist or non-literal interpretation of moral language.

Many philosophers, however, want neither of these options. They want to say that they find at least some substantive moral judgements, literally understood, to be true. Chattel slavery is unjust. The infliction of pain upon others for one’s own amusement is wrong.<sup>3</sup> Yet some of these same philosophers want to counten-

<sup>1</sup> I am indebted here to Gideon Rosen. See P. Railton and G. Rosen, ‘Realism’, in J. Kim and E. Sosa (eds.), *The Blackwell Companion to Metaphysics* (Oxford: Blackwell, 1994).

<sup>2</sup> The best-known example of an error theory in the moral case is that of J. L. Mackie, *Ethics: Inventing Right and Wrong* (London: Penguin, 1977).

<sup>3</sup> It is not important here whether we understand these judgements as absolute or *prima facie*. Even a moral realist can, for example, espouse a theory of *prima facie* moral reasons, or a

ance moral facts only so far – in particular, they wish to distance themselves from moral realism.

One way to avoid moral metaphysics while accepting the idea that at least some moral judgements, literally interpreted, are true would be to adopt non-literalism about truth itself (at least as it figures in moral discourse). But this option has already been foreclosed in our present discussion, since we are supposing that when the term 'truth' is employed in moral discourse it has its familiar sense, the same sense in which it figures in discourse about those parts of the world toward which we more readily adopt a realist attitude. Moreover, if there is a difference of metaphysical interest between moral facts and facts about garden-variety objects, we have been assuming, this is not because 'fact' has a different meaning in the two cases.<sup>4</sup>

Let us call *Literal Truth* the position that at least some substantive moral statements, taken literally, are literally true. What space remains between Literal Truth and moral realism?<sup>5</sup> Perhaps the space is only rhetorical – there is nothing a moral realist could legitimately want beyond what is contained in Literal Truth. That would be reminiscent of a familiar realist/anti-realist dialectic in which the anti-realist places the burden upon the realist to say what more could be at stake. But in this case the dialectic is different. For here it seems as if the party with the stronger interest in showing there to be something genuine at issue is not the realist (who accepts Literal Truth as a matter of course) but the moral-anti-realist-who-nonetheless-accepts-Literal-Truth. This sort of anti-realist wants to be able to distinguish morality from areas of inquiry that he thinks to be more properly informative about the shape of the real world.

A qualification is in order. Given the (uncontroversial) super-venience of moral discourse, to insist that at least some moral judgements, literally interpreted, are true is to commit oneself to

---

theory that involves vagueness and indeterminacy. We cannot saddle the moral realist with standards of absoluteness or 'rational determinacy' we do not impose upon such 'hard' domains of inquiry as biology or geology. (Just which genetic change made for the first member of *Mammalia*?) Indeed, it is commonly argued in philosophy of science that the choice of a fundamental physical theory is not 'rationally determinate', even given ideal evidence.

<sup>4</sup> Obviously, it would not help at this point to embrace a global anti-realism about factual discourse, since precisely the idea is to work out when a realist attitude toward a factual domain of discourse is sensible and when not.

<sup>5</sup> None, according to some commentators. See for example G. Sayre-McCord's introduction to his *Essays on Moral Realism* (Ithaca: Cornell University Press, 1989).

the existence in the real world – where else? – of some or other non-moral states of affairs in virtue of which the moral judgements hold. So Literal Truth by itself commits one to the idea that moral discourse reveals something of the shape of the real world after all: the world must be such as to underwrite at least some bona fide moral facts.

That could be very metaphysically informative – or not – depending upon the substantive content of moral claims, a question that cannot be bracketed in discussions of the significance of Literal Truth and that we ourselves cannot any longer postpone. One potential attraction of a position that combines Literal Truth with anti-realism about morality is that it might provide a way of capturing what has seemed to many right about emotivism while avoiding the imputation of a non-standard grammar to moral discourse and the attendant exaggerated attention to non-cognitive states. Emotivists explained the fact/value distinction by arguing that moral discourse has a grammatically *subjective* character – unlike ordinary objectual discourse it functions primarily to express the speaker's emotions rather than to describe states of affairs. Philosophers who combine Literal Truth with anti-realism will reject this grammatical way of accommodating the subjective aspect of moral discourse and the associated view of value claims as non-factual. But they can adopt the emotivists' re-orientation of moral philosophy away from the idea that moral judgement passively reports objectual states worthy of reification and toward the idea that it essentially involves the reactions and actions of subjects. This re-orientation should inhibit reification and realism, even if one grants Literal Truth. By contrast, various judgements purporting to describe the external world, other minds, and the past can be made seemingly without any necessary engagement on our part as subjects, which explains why reification and realism come much more naturally to them in the wake of Literal Truth.

Consider in this light what two philosophers have said recently about what realism would appear to involve in the moral case. John Rawls characterizes the realist's conception of moral inquiry as:

the search for moral truth interpreted as fixed by a prior and independent order of objects and relations, whether natural or divine, an order apart and distinct from how we conceive ourselves.<sup>6</sup>

<sup>6</sup> Rawls, 'Kantian Constructivism in Moral Theory', *Journal of Philosophy* 77 (1980), p. 519. Note also the discussion in his more recent book, *Political Liberalism* (New York: Columbia University Press, 1993), lecture 3.

And in a similar vein, Crispin Wright writes in concluding a discussion of moral realism using a comparison with realism in science:

[according to the realist] the real progress of science is measured by the extent to which our theories represent a reality whose nature owes nothing to our natures or the standards that inform our conception of reasonable discourse about it.<sup>7</sup>

If a realist conception of moral facts must be one in which they are 'prior and independent', or 'a reality whose nature owes nothing to our natures', then it would be clear why space remains between Literal Truth and moral realism: an anti-realist could claim that moral facts lack the metaphysical independence from human subjectivity that would sustain a realist attitude toward them. This would not be a defect in the moral facts themselves, but rather a ground for recognising that realist reification is not always licensed by a willingness to speak literally of truth and fact.

The question now becomes: Does a realist attitude necessarily bring with it a commitment to this sort of metaphysical independence?

## II

Certainly some forms of realism do owe their distinctive character to this sort of commitment. Such paradigms of realism as realism about the external world, realism about the theoretical entities of natural science, and Platonism about mathematical objects have as a central component the idea that the entities or properties posited are 'prior to and independent of' any condition or experience of subjects. In that sense, they are radically non-subjective. Since 'subjective' is a notoriously slippery term, let us coin the technical term 'subject-ive' (with a hyphen) to express the notion of that which is essentially connected with the existence or experiences of subjects, i.e., beings possessing minds and points of view, being capable of forming thoughts and intentions. These paradigms of realism are, in our hiccoughing neologism, radically nonsubject-ive.

Similarly, many paradigm forms of *anti*-realism take the subject-ive as their point of departure: phenomenologists interpret physical object language in terms of actual or possible observations; observationalists do the same with the theoretical language of

<sup>7</sup> Wright, *Truth and Objectivity* (Cambridge: Harvard University Press, 1992), p. 200.

science; conventionalists trace the nature of numbers as well as the force of the logical 'must' to the practices of the mathematical community; nominalists see universals as linguistic rather than metaphysical; and atheists treat God as a mere idealisation of human characteristics ('The secret of the Holy Family is the human family', according to the Feuerbachian). It should be emphasised, however, that subject-ive does not mean subjective, at least insofar as that term is used derogatorily to suggest a domain without standards, where arbitrary opinion takes the place of judgement. Phenomenalists and observationalists have characteristically emphasised such standards of epistemic objectivity as intersubjective agreement and reliance upon rational credence functions. What these views reject in realism is not its interest in objectivity, but (what we might call) its tendency toward objectivity – its reification of a domain as independent of subject-ivity.

Radical nonsubject-ivity, where it can sensibly be sustained, does help to secure several seeming hallmarks of realist thought. If a domain of entities and properties *D* is independent of experience, for example, then it has at least three interesting kinds of explanatory potential. First, it can in principle support non-circular explanations of the patterns in our experience – we see a tree in the quad because one is there, and not vice versa. This corresponds to a certain notion of univocal 'order of explanation'.<sup>8</sup> Second, *D* can in principle support explanations unmediated by any sort of experience – a tree falling unnoticed in the quad nonetheless leaves an elongated depression in the earth. This, taken together with the first, establishes an impressive 'width of cosmological role' for *D*.<sup>9</sup> And third, these two kinds of explanatory potential enable us to see clearly how explanations invoking the properties and entities of *D* can in principle pass an 'Attitude Test' of a kind suggested by Gilbert Harman: if we were to replace reference to these properties and entities in our explanations with reference only to our attitudes about *D* combined with associated non-*D* facts, genuine explanatory information would be lost.<sup>10</sup>

<sup>8</sup> Compare the discussion of 'order of determination' and its relation to realism, in the essays by M. Johnston and C. Wright in J. Haldane and C. Wright (eds.), *Reality, Representation and Projection* (Oxford: Oxford University Press, 1992).

<sup>9</sup> This useful term is due to Crispin Wright. See his *Truth and Objectivity*, ch. 5.

<sup>10</sup> See the discussion in ch. 1 of G. Harman, *The Nature of Morality* (New York: Oxford University Press, 1977). The notion of 'explanatory information' is roughly this: information about what a comprehensive explanatory history would contain. For discussion, see P. Railton, 'Probability, Explanation, and Information', *Synthese* 48 (1981).

(Contrast the atheist's view of religious explanations: world history has been affected in countless ways by religious belief and practice, and religious conviction has led many to have experiences which they call miracles or sensations of the presence of God, but no actual deity has ever caused any perception or even the swerving of a single atom. For the atheist, then, if religious explanations were replaced systematically with explanation in terms of religious attitudes and collateral non-religious phenomena, no genuine explanation would be left out. The theistic realist, by contrast, would insist that in many cases no such explanatory replacement is possible.) These three notions of explanatory potential give us an idea of the particular sort of explanatory pay-off that positing a domain of nonsubject-ive entities and properties can yield, and which might well make such a posit worthwhile.

It is natural to associate realism with this sort of radical nonsubject-ivity. The realist is said to want entities and properties that are 'out there' in the world, independent of the projections of our beliefs and indifferent to the tossings and turnings of our imaginations – entities and properties that are there *anyway*.<sup>11</sup> This association bodes ill for the moral case. We have a degree of confidence – perhaps even justified confidence? – that we have some idea of what it would be like for rocks and trees, quasars and molecules, heat and mean molecular kinetic energy to be 'prior to and independent of' all subject-ivity. By contrast, the very idea that there might be a realm of moral properties subsisting independent of all subject-ive phenomena – all facts about human nature or experience, say – would be philosophical adventurism of the first order. Thus interpreted, moral realism is a non-starter.

### III

But is this train of thought compelling? It overlooks the seeming platitude that the nature of realism about any domain depends upon the nature of that domain itself. What do we imagine the entities or properties of this domain to be like? If there were such things, what would they do? The question of realism about a domain *D* of purported entities and properties cannot be treated in a purely generic way – we must ask more specifically what would it take to fill *D*'s bill.

Let us try to make this platitudinous thought more explicit.

<sup>11</sup> I owe this phrase to Gideon Rosen.

Suppose we were, for a given domain *D*, to draw up a 'job description' for purported *D*-entities or *D*-properties – '*D*'s' for short – based upon our ordinary notions, folk theories as influenced by scientific developments, relatively uncontroversial applications, paradigm cases, and so on.<sup>12</sup> This job description is to be very literal-minded, that is, it attempts to express what *D*'s *if literally understood* are supposed to be and do.

For example, a job description for *physical object* would include not only such minimal conditions as spatio-temporal location and continuity, but also all manner of humdrum activities and folk-scientific functions. Thus physical objects persist when unobserved; they non-circularly explain various features of our sensory experience; they have properties which enjoy a certain independence from our conceptions of them; and they include, as paradigms, rocks and trees; and so on. The presence in our intuitive job description for physical objects of these various elements is reflected in the fact that almost everyone views phenomenalism as a *revisionary* or non-literal account of our physical object discourse.

Job descriptions therefore are hardly free of folk-theorising. As we noted earlier, mental states are commonsensically supposed to be internal conditions of sentient beings that non-circularly explain their manifest behaviour. Analytic behaviourism thus is widely seen as a revisionary account of our mentalistic discourse. The Christian God is supposed to be a person-like entity with

<sup>12</sup> This way of approaching things is obviously inspired by the Ramsey-Carnap-Lewis approach to defining theoretical terms. (See, e.g., D. Lewis, 'How to Define Theoretical Terms', reprinted in his *Philosophical Papers*, vol. 1 [New York: Oxford University Press, 1989].) Unlike the classic forms of this approach, however, there is no assumption here that the job description will be 'observational' in the sense of being innocent of 'theoretical' predicate constants. Moreover, as in Lewis's more recent versions of this approach, it is assumed that the job description is modified by a qualifier to the effect that entities or properties can meet the description more or less fully.

Lewis himself does not offer this sort of account of moral terms (see his 'Dispositional Theories of Value', *Proceedings of the Aristotelian Society* suppl. vol. 63 [1989]), preferring a more explicit definition. Defending a Ramsey-Carnap-Lewis approach to a class of terms requires that the corresponding properties purport to have the right sort and extent of theoretical roles. Certainly this can be doubted in the moral case. I would argue that moral discourse is well-integrated into our folk explanatory theories, which include counterfactual-supporting generalisations linking needs and wants, goods and functions, virtues and stable personality traits, and rightness and social functionality (see for example the discussion in P. Railton, 'Moral Realism', *Philosophical Review* 87 (1986) and 'Naturalism and Prescriptivity', *Social Philosophy and Policy* 7 (1989)). These theories attempt to explain, among other things, both why we sometimes arrive at moral knowledge or act well and why we sometimes do not.



extraordinary powers. An interpretation of 'God exists' as 'I hereby commit myself to loving my neighbour' would, I suspect, strike most theists as non-literal.

Some elements of job descriptions are more epistemic in character. In the case of mental states, we take ourselves to have a kind of access to at least some of our own mental states that is more direct than our access to the mental states of others. In the case of physical objects, we take ourselves to know something of them through experience. In the case of numbers, too, we take ourselves to know something of them, but with a special kind of certainty, and through means seemingly quite unlike those of ordinary observation. A radically empiricist interpretation of mathematics, like Mill's, is to some degree epistemically revisionary.

But to be revisionary is not to be wrong. Ordinary discourse is bound to contain not only vagueness and ambiguity, but inconsistencies, incoherences, impossibilities. Ordinary discourse shows the sedimentation and metamorphosis of centuries of folk, scientific, and quasi-scientific thought, and any account free of ambiguity and incoherence is bound to be somewhat revisionary. Job descriptions are literal, but it seldom is possible that they be literally met when fully spelled out. They are meant to map out explicitly the intuitive landscape toward which philosophical accounts are directed, alternatively claiming with pride to capture large tracts of terrain or arguing defensively that those parts they cannot fit are best got rid of.

Any such attempt at explicitness about what is uncontroversial is itself of course bound to be somewhat controversial, especially around the edges. Just *how* folk – as opposed to scientifically informed – a conception do we follow? And just *which* folk do we consult? Common sense is not regimented or homogenous. Consider 'freedom of the will'. It certainly appears to be part of our commonsense conception of free will that free agents could have acted otherwise. But is it part of this conception – or a philosophical invention meant to explain it – that agents are able to act contra-causally? Moreover, in some languages the commonsense term is closer to 'free choice' than 'free will'. Does our folk conception of free action involve commitment to the will as a little piece of machinery turning the psychic gears inside the agent? Or would a belief-desire-based model of choice more typical of contemporary psychology suffice?<sup>13</sup> It will always be an essential

<sup>13</sup> I am grateful here to Jean-Pierre Dupuy.

part of revisionist strategy to claim that what seems to some to be part of the intuitive landscape is really a mirage or a philosophical excrescence.

#### IV

With these caveats about 'reading off' job descriptions from common sense or expecting them to be fully met in mind, let us ask what the job listing for moral properties might look like. This is a fairly broad terrain, including properties as diverse as those of obligation, value, and character. To give the full listing would involve a large number of more or less interdependent job descriptions. Moreover, moral properties are asked to do a lot of things. We cite them in deliberation, assessment, interpretation, instruction, punishment, prediction, and explanation. Their job descriptions would be complex, and much intertwined with other elements of commonsense thought. A moral realist need not be non-revisionary with respect to the whole lot. And a moral realist need not be a realist about the whole lot. Just as one can be a scientific realist about entities and events but not about probability or nomic necessity, one can be a moral realist about value and virtue but not duty and rights.

Let us focus herein on questions of moral value and evaluation (rather than claims about moral duty, say), and pay special attention initially to some questions about explanation that have figured centrally in recent discussions of moral realism. And let us remember our larger purpose. Literal Truth about a domain *D* is committed to the idea that the job descriptions of enough central *D* entities or properties are sufficiently well met by conditions in the actual world to enable us to formulate some substantive truths about *D*'s as literally understood. Realism in turn involves Literal Truth. Does it necessarily involve a further demand that the entities and properties of *D* be of an independent or nonsubject-ive nature?

Consider first the realist about the external world of physical objects. Our ordinary conception of physical objects has it that they affect not only consciously-formulated beliefs ('I think we've struck a submerged reef') but also experiential states of a kind prior to conceptualised belief ('We test to see if the newborn can follow a flashlight with his eyes'), non-experiential states of subjective beings ('I always get freckles like this in the summer sun'), and radically nonsubject-ive states ('The earth's gravitational field

holds the moon in its orbit').<sup>14</sup> As we noted earlier, it is all part of a day's work for physical objects to have effects in the world unmediated by and independent of any conscious or subject-ive state. An interpretation of physical object discourse bold enough to deny this would be bound to be viewed as too revisionary to vindicate our ordinary notion.

For this reason, nonsubject-ivity belongs to the job description of physical objects, and therefore is as much a commitment of Literal Truth about physical object discourse as it is an implication of realism about physical objects. To believe that at least some physical object statements, literally understood, are literally true, is not to accept any particular explanatory claims (except perhaps, with the usual qualifications, for paradigm cases). But it is to accept the existence of entities with this sort of explanatory potential as part of the furnishings of the actual world.

Contrast the case of belief. Common sense has it that beliefs are the possessions of subjects, and never found in their absence. One would not know what to make of a demand that beliefs as such, as ordinarily understood, have an explanatory role in the cosmos independent of all consciousness and unmediated by thought or its embodiment. The chief explanatory roles figuring in the job description of belief concern the shaping of thought, inference, action, perception, expectation, emotion, dreams, and so on. The effects of beliefs on the wider, nonsubject-ive world are always mediated by these states and processes and their embodiment in subject-ive beings. To be sure, their embodiment itself can have unmediated effects on the nonsubject-ive world. The pattern of electrical activation of my neurons as I ponder how to turn a key in a stubborn lock subtly affects the electromagnetic field surrounding me. But beliefs 'in their own right' – beliefs as such rather than their embodying conditions – effect changes in the physical world only indirectly, thanks to their influential role in the lives of the subjects wherein they reside.

A requirement that beliefs constitute a 'prior and independent order' or possess a 'wide cosmological role' would thus certainly not be part the conditions for Literal Truth about belief discourse. Might it be an additional condition necessary to move from Literal Truth to realism? Could a non-eccentric realist about belief afford to be so eccentrically nonsubject-ive about the nature of belief? In

<sup>14</sup> This brief list follows Wright's description of 'wide cosmological role' in *Truth and Objectivity*, pp. 196–99.

principle, a realist can espouse a tolerable revisionism, bending our ordinary conceptions in order to obtain a good enough fit overall. Yet it is hard to see what positive contribution would be made to the vindication of commonsense or scientific belief discourse by attempting to introduce beliefs without believers.

More generally, it need be no part of the ambition of realism about a domain of discourse *D* to attribute to the entities or properties in *D* causal-explanatory roles or conditions that go beyond the roles and paradigm cases found among the 'job listings' of *D*. Of course, any given realist might insist upon such further roles in order to meet other philosophical desiderata. A naturalist might insist as a condition for a realist attitude toward *D* that a naturalistic reduction of *D*'s be in the offing, or that *D*'s win their way into our best scientific theories rather than remain at the level of commonsense explanatory frameworks. But this naturalistic condition would be a specific way of being a realist or motivating realism, not a requisite of realism in general. And no matter what stripe of realism is in question, it will tend to complicate the realist's life to insist upon a role or standing for *D*'s that goes against the grain of the job listings for *D*. A starkly object-ivist view of so paradigmatically subject-ive a phenomenon as belief might more appropriately be viewed as a replacement theory rather than a form of realism. Similarly, the contemporary scientist who says that God is, for him, simply the laws of physics themselves is more likely to be seen as offering a substitute for traditional notions of God than a vindication of them.

If we are to make sense of the question of realism about belief, then there must be a way of being a realist about subject-ive domains. But what of the notion that realism about a domain *D* is distinguished by its commitment to the 'objective' status of *D*'s? Is there nothing to this? Or, could there be so little to realism?

Beliefs are fundamentally subject-ive, given their jobs, but it must be kept in mind that the sense of 'subject-ive' is technical and stipulative. It is opposed not to 'objective', but to 'object-ive', that is, having to do essentially with nonsubject-ive entities – entities that lack a mind or point of view, entities that are not a locus of experience or intention. Knowledge, for example, is by its nature subject-ive since it involves belief. But though it therefore cannot be object-ive, it might well be objective. And so might belief and belief-attribution. Consider the 'Attitude Test' mentioned earlier.<sup>15</sup>

<sup>15</sup> There are numerous dimensions to our none-too-well-understood notion of objectivity. Here we will be concerned mostly with those aspects reflected in the Attitude Test. For a

Can we uniformly replace the purported explanatory role of beliefs by reference instead to what we take ourselves to believe, without any loss of explanation? Or does belief have the potential to support explanations even in the face of what we take ourselves to believe? According to a familiar strain of anti-realism about belief, there are no facts about what we believe which obtain 'outside' an interpretive scheme, so that the whole contribution of belief-talk to understanding human behaviour can be captured by talking instead of belief-attributions, norms of belief-attribution, and so on. Realists about belief, on the contrary, think that there are facts about what an individual believes that can contribute to the explanation of his behaviour in ways not dependent upon his interpretive scheme or ours.

Realists here take their cue in part from commonsense thought about belief. We earlier noted that belief is commonsensically treated as an internal state of individuals that has a causal or quasi-causal role in shaping their conduct, sometimes consciously sometimes not. If common sense is right about belief, then belief attributions and interpretive schemes are answerable to facts about the distribution of beliefs, not the other way around. Belief discourse would pass the Attitude Test and, to that extent, possess a kind of objectivity. Because its explanatory role is always mediated by subjects, belief has narrow explanatory scope; but because this role is not always mediated by self-conception or interpretation, belief can be an objective feature of the subject-ive part of the world. Philosophical approaches to belief that treat belief as radically indeterminate and interpretation-dependent deny that belief possesses this sort of objectivity, and therefore strike the man in the street as quite a novelty (at least in my experience), surprisingly at odds with a literal understanding what beliefs are supposed to be. Indeed, such approaches collide with our ordinary way of speaking of ourselves in profound ways, and even flirt with incoherence.<sup>16</sup>

What, then of the moral case? If moral value were a something rather than a nothing, what sort of a something would it be? And

---

fuller description of some notions of objectivity, see P. Railton, 'Marxism and Scientific Objectivity', reprinted in R. Boyd, P. Gasper, and J. D. Trout (eds.), *The Philosophy of Science* (Cambridge: MIT Press, 1991).

<sup>16</sup> For relevant discussion of the possible incoherence of another sort of denial of Literal Truth about belief discourse, see P. Boghossian, 'The Status of Content', *Philosophical Review* 99 (1989).

what would it do, or help to explain? What would its paradigm cases be? Although it is no ambition of this paper to sketch out a complete idea of the job description of any moral property, we do need to ask: Would moral value be subject-ive, like belief? And, if so, would it be capable of (at least) the same sort of objectivity as belief?

It certainly seems safe to say that moral value has essentially to do with subjects. Just as a world altogether without conscious beings would contain no beliefs, so would it lack moral phenomena. The development of moral philosophy throughout the modern period has been deeply influenced by Hume's observation of a special connection between moral evaluation and action, and by Kant's insistence that moral reasoning must be practical. The exact nature and modality of the connections between moral thought and the experience or agency of subjects are matters of continuing controversy within moral philosophy. But it is relatively uncontroversial that commonsense moral discourse supposes there to be some such connections – e.g., that 'ought' implies 'can' and that evaluation is normally associated with motivation – and that any satisfactory account of morality must somehow accommodate or otherwise explain these ties. We have numerous paradigm cases of moral value or rightness, but so far as I know none that concern wholly nonsubject-ive phenomena.<sup>17</sup>

Since the job description is subject-ive in this way, there can be no question of Literal Truth or realism about moral properties as requiring a 'prior and independent order'. Yet, according to either realism or Literal Truth, moral properties might nonetheless possess the sort of objectivity we found in the case of belief.

Consider some examples of commonsense moral explanations. The injustice of segregation helps explain why it produced widespread alienation and discontent and an eventual movement for change. The honest decency of a colleague helps explain why she has come to be listened to with some care in controversial matters. The deviousness of a parent helps explain the insecurity of his child later in life. And so on. Such explanations are rough but informative. For example, we learn something when an observer attributes an important role in explaining a country's repeated

<sup>17</sup> This is not true of value discourse in general. Even if aesthetic value must in the end be understood as somehow a relational matter involving subjects, still, many paradigm cases of objects possessing aesthetic value are inanimate objects.

bouts of social instability to underlying injustices rather than to economic downturns or political factors alone.

Such moral explanations – at least, when offered non-metaphorically – have narrow ambitions regarding cosmic scope. Their explanada are always in the first instance features of subjects.<sup>18</sup> Subjects can of course go on in various ways to change the nonsubject-ive world – popular rebellions attributable in part to social injustice have, for example, led to changes in settlement patterns or agricultural practices with profound ecological effects. What matters most for our purposes, however, is that many moral explanations would pass the Attitude Test in the following sense: they could not be replaced without loss by an explanation adverting only to the moral beliefs of those involved. A social order's injustice may produce alienation and dysfunction well before any articulated sense of its injustice surfaces within the population. A parent's dishonesty may undermine the solidity of a child's sense of self long before any glimmer of a moral critique of the parent emerges in the child's mind – if indeed it ever does. These explanatory accounts are in the first instance unmediated by moral concepts or beliefs. It would be quite a different analysis of a situation to trace the origins of a child's difficulties to the acceptance of a morally-conceptualised *belief* on his part that his father is behaving deceptively.

Moral sceptics and nonfactualists will on general grounds refuse to take any moral explanation at face value, for they deny the existence of moral properties. If these purported explanations are at all informative, they will argue, it is because of what such accounts tell us indirectly about moral beliefs, accepted norms, and other non-moral conditions. But here we are addressing the philosopher who has already accepted Literal Truth in the moral case. And commonsense moral discourse, taken literally, is so thoroughly interwoven with commonsense theorising about human action, society, and history, that it seems an evident part of the full job description of moral properties for them to have various genuine explanatory roles.

No full job description for moral virtue, for example, could leave it explanatorily inert. It is of the essence of virtue that

<sup>18</sup> We set aside here, as we did in the case of belief, explanations couched entirely in terms of the conditions of *embodiment* by subjects of the properties in question. Clearly the 'subjacent base' of moral properties, as R. M. Hare has called it, can directly explain features of the world without subject-ive mediation.

possession of a virtue would help explain constancies of individual behaviour across varying circumstances, and that variations in behaviour across individuals facing similar circumstances can be attributed to the presence or absence of virtues.<sup>19</sup> Moreover, it is of the essence of virtue-based explanations that they are *objective* in the sense discussed above – it is one's underlying moral character, not simply one's moral beliefs (and certainly not one's beliefs about one's character), that explains one's conduct. The question of whether to accept certain moral explanations as objective therefore need not derive from concerns or reifying tendencies specific to realism at all. It is difficult to imagine acceptance of the literal truth of substantive claims involving moral virtues without acceptance of the idea that virtues afford some measure of objective explanation (at least to the extent of passing the Attitude Test).

If differences over how fully moral properties can do their purported jobs – including their purported objective explanatory jobs – belong to the space of revisionism within Literal Truth as much as they do to realism, what sorts of differences might distinguish Literal Truth from moral realism? We have already seen that they cannot be differences over the subject-ive character of moral properties – that feature of moral job descriptions is sufficiently central to be common ground to all interpretations of moral discourse with any hope of avoiding error theory. Could it be that 'realist' is simply the self-description of choice of those who think that the job descriptions of moral properties can fairly fully be met – so that adopting the position Literal Truth while rejecting the label 'moral realism' would be a way of signalling that one thinks some substantial, but still tolerable, revisionism of common sense is called for? There is something to this thought, but not, I suspect, enough. More than likely, there will be self-described moral realists who are more revisionary with respect to ordinary moral discourse than some self-styled quietists who accept Literal Truth but balk at moral realism.

Differences about the appropriateness of an attitude of realism toward the moral among those accepting Literal Truth will, I suspect, lie less in matters about whether to accept at face value

<sup>19</sup> These features of discourse about moral virtue testify both the empirical content and the empirical vulnerability of moral explanation. At least one important tendency in contemporary psychology is sceptical about the existence of continuing traits of character capable of explaining patterns of constancies and variability in individual conduct across situations. See L. Ross and R. Nisbett, *The Person and the Situation* (Philadelphia: Temple University Press, 1994), esp. ch. 4.



moral explanations than in general interpretive questions about the metaphysical weight of terms like 'cause', 'explain', 'property', and 'objective'. In any event, our concern about whether the job descriptions of moral properties – including the explanatory elements – can be met does not itself depend upon any 'heavy-weight' notion of 'explain' or 'property'. What is at issue here is a matter common to Literal Truth and realism: whether a large constellation of platitudes, commonsensical practices, folk notions, imported scientific theorising, and paradigm cases can – to a tolerable degree – be found to be instantiated in this world.

## V

Given the differences in job listings between the domain of physical objects and the moral domain, it is unsurprising that moral judgement is a different business from ordinary judgement about objects in the world around us. Moral properties could hardly do their job if they were radically object-ive, physical objects could hardly do their job if they were not.

Belief furnishes us a model of how there can be a genuine issue of realism and objectivity concerning a subject-ive domain. But in emphasising the parallel between moral realism and realism about belief, I do not mean to suggest that moral realism must be a form of psychologism. For commonsense moral thought includes various roles for moral properties that are not obviously psychological. Most notably, there are the so-called *normative* roles of moral notions – the action-guiding character of terms like 'ought', 'must', and 'good'. It is an open question whether any purely psychological theory can give an account of this normativity.

Yet here, too, an interesting parallel with belief holds. For it is also an open question whether any purely psychological theory can give an account of belief. Beliefs are commonsensically attributed propositional content as well as causal roles. Moreover, these contents have correctness conditions that are commonsensically seen as objective in the sense that the content of one's beliefs is not simply a matter of what one takes it to be – it is not a matter of idiosyncratic will, free stipulation, or spontaneous creation. If the job description of belief is to be met in its central elements, an essentially subject-ive state must be capable of possessing the sort of objective normativity embodied in correctness conditions. But the puzzle here does not arise from the fact that the state is subject-ive as such. It is no more clear – on the contrary, it is a good deal

less clear – how a purely nonsubject-ive state of the world could be a bearer of content. If the feat of creating and embodying states with correctness conditions can be accomplished at all, it seems evident that subjects rather than mere objects will be the ones to do it.

The puzzle is *how* they might do this: how subjects might, by doing what they do, place themselves within a normative framework that sustains a distinction between what is correct and what is done (including what is as a matter of fact taken to be correct). This is as much a puzzle in the case of belief as it is in the moral case, though not the same puzzle – morality is certainly to that extent different. And it is a puzzle for Literal Truth even apart from realism. For if moral discourse or discourse about belief is to any significant degree literally true, then somehow we subjects must have brought objective normativity into being.<sup>20</sup>

*Philosophy Department*  
*University of Michigan*  
*Ann Arbor, MI 48109–1432*  
*USA*

<sup>20</sup> I would like to thank Paul Boghossian, Allan Gibbard, Paul Horwich, David Lewis, Gideon Rosen, Michael Smith, and Crispin Wright for much helpful conversation. An earlier version of this paper was presented at an NEH Summer Seminar at the University of Nebraska and at the Centre de Recherche en Epistémologie Appliquée, Ecole Polytechnique; I am grateful to those attending for many useful comments and criticisms.