PLURALISM

The philosophy and politics of diversity

Edited by
Maria Baghramian
and
Attracta Ingram

London and New York

3

A
SENSIBLE PERSPECTIVISM

Philip Pettit

Perspectivism is associated with the work of Nietzsche in the nineteenth century and it is frequently invoked nowadays as a theme in postmodernist writings; in particular as a ground that justifies radical forms of relativism and irrationalism. There is no prospect of rational reconciliation between divergent viewpoints, under the picture that perspectivism allegedly delivers. Anything goes in principle, and what goes in practice is a function of where the power and the influence lie.

The core perspectivist idea is that all claims to knowledge bear the marks of the contingent circumstances and assumptions of those making the claims, so that it is deeply mistaken to privilege any particular perspective and treat it as providing the ultimate heuristic standpoint. I find the core perspectivist idea attractive but I do not think that it leads towards the iconoclastic position with which it is associated. This essay is an attempt to identify what is right in the idea and to provide an overall framework for making sense of it.

The essay does not look at the political implications of the sort of perspectivism defended, but these are hardly going to be iconoclastic either. The moderate perspectivism envisaged fits naturally with a pluralist, tolerant attitude towards divergence in modes of inquiry and evaluation – indeed it suggests that any other attitude would be seriously out of place – and for that reason it encourages the pluralist vision of a society and polity where different viewpoints can be fostered. But it does not associate such pluralist sentiments with the rampant irrationalism that usually goes under perspectivist colours.

The paper is in five sections. First I present a simple case for perspectivism, explaining what exactly I mean by the doctrine. Then I chart three possible responses to it, only one of which looks attractive: the belief that perspectives can be integratively reduced. I provide a model of integrative reduction – the reduction of the indexical to the non-indexical – in the third section; and then in the fourth I try to make plausible the claim that this model can be extended to make sense of discursive, perspectival cleavages in general. The paper concludes with a section in which I discuss the significance of perspectivism, as it is understood and accommodated here, for philosophy in general; I argue that it gives a new slant on what we should expect philosophy to try, and to be able, to accomplish.

A case for perspectivism

The core perspectivist idea is familiar to contemporary analytical philosophers from Thomas Nagel's claim that there is no such thing as the view from nowhere (Nagel 1986). All representations of what there is are affected in one way or another by contingencies, roughly, of position. They are marked by the sensory organisation, cognitive processing or linguistic articulation on which they rely, by the theoretical assumptions and cultural associations that they embody, or by some such contingently variable influences. All forms of representation are perspectival and there is no hope of transcending those perspectives in a god's-eye depiction of things (see too McDowell 1983; Nagel 1986; Price 1992; Strawson 1985; Williams 1978).

I am happy to go along with this story of multiple viewpoints or perspectives. The reason is that a relatively straightforward line of argument seems to me to support the position. I will sketch that argument briefly in this section, drawing on work that I have done elsewhere (Pettit 1990, 1991, 1994). The argument is Wittgensteinian in character but I do not try to detail or establish its Wittgensteinian credentials.

The argument can be presented in seven steps.

1 Any discursive form of knowledge presupposes semantic competence in the terms of the discourse.

2 Some of the terms within any subject's competence must be semantically basic: that is, understood ostensively, without theoretical definition.

3 No finite set of ostended exemplars instantiates just one potential semantic value – just one referent – for a term.

4 If such exemplars identify a relatively determinate value for a subject, therefore, that must be because of a contingency in how the subject is disposed to respond.

5 Thus the mastery of the basic terms in any subject's vocabulary is achieved in a response-dependent way; the unresponsive subject will lack the conditions required for non-parasitically mastering such a term or concept.

6 There are different discourses to the extent that the responsiveness required for semantic competence is tuned or primed by different clustered sets of factors.

7 Perspectivism is best understood as nothing more or less than the claim that there are different discourses in this sense.

Steps 1 and 2

The argument begins with two simple but irresistible steps: first, the observation that I can be possessed of a certain form of discursive knowledge only so far as I grasp the meanings and the referents of the terms in question; and second the remark that I must understand some of those terms in a basic, theoretically undefined way. Many terms will be understood on the basis of definitions, explicit or implicit, in the other terms at my disposal. But, on pain of circularity, not all terms can be understood in this way. Some of them I must master just from seeing how they are used: in particular, what they are used to pick out. I must master them on the basis of ostension, or something like ostension: that is, by having sample referents of the terms brought to my attention.

While these claims are hard to resist, it is worth adding two caveats. First, the terms that I master ostensively may not be mastered one by one, in an atomistic fashion; it may be that I can only master a number of such terms together in the sort of package deal that Quineans talk of; it may even be that I can only master some terms ostensively so far as I am already the master of certain other terms, ostensive and defined. Second, the terms that I master ostensively at any one point in time may come later to be mastered definitionally; and the terms that I master ostensively may be understood definitionally by others, and vice versa. The only claim here is that at any particular time each of us must understand some of the terms within our competence on a pretheoretical, ostensive basis.

Step 3

The third step in the argument makes the familiar point that no matter how many instances are presented to me as ostended examples of where a given term applies, those instances will fail to determine a unique semantic value for the term (Kripke 1982). Let the term in question be a predicate, so that the semantic value it is to be assigned is a property. No matter how many instances of the property are presented as examples, they will not succeed in uniquely determining that property as the semantic value ostended. The reason is that any finite set of things will instantiate an infinite number of properties: it will be a proper sub-set of an infinite number of possible property-extensions.

Suppose I am presented, for instance, with a set of regular geometric shapes. Those examples will certainly instantiate the property of regularity but they will also instantiate an endless variety of other, perhaps gerrymandered properties: properties like that of being regular-in-the-shapes-presented-or-otherwise-being-irregular, of being regular-both-in-the-shapes-presented-and-in-all-three-sided-shapes-or-otherwise-being-irregular, and so on through an open-ended set of possible variations.

Steps 4 and 5

The fourth and fifth steps draw the lesson taught by the last. We know that human beings are capable of learning the semantic value of a basic term – learning which property is to be associated, for example, with a basic predicate – from ostension. And that means that a finite set of instances can make a property salient to people and teach them which property to associate with a given predicate. Since this trick cannot be pulled off just on the basis of the instances, it must involve something happening on the side of the subject, as a result of exposure to those instances. It must be that, presented with those examples, the subject responds with an inclination to go one way rather than another in extrapolating from the set of examples and identifying things that share, and things that do not share, the ostended property. The ability to master the term and concept in question – the ability to learn which property it designates – is response-dependent. It presupposes that the subject is tuned to respond in a certain extrapolative way and thereby to latch onto the particular property associated with the extrapolative disposition. Otherwise the set of examples would leave the subject at sea, aware of too many candidates for the role of providing the semantic value of the predicate.

As we needed to enter caveats before, so it is important to mention a number at this point too. First, the response-dependent story of semantic mastery that I have sketched is not meant to be a complete account of conceptual mastery or, more generally, rule-following. At the least, it needs to be filled out by a story as to how a subject can get a property wrong – and be conscious of this possibility – if they rely on their own extrapolative inclination to identify instances (Pettit 1990, 1996).

Second caveat. The response-dependent story – say, as filled out to allow for fallibility – does not aspire to show that ostension makes a property determinate to the point of fixing the right predication in all possible novel cases (*pace* Kripke 1982; see Pettit 1996, Postscript). What it is introduced to vindicate is the capacity of ostension to fix the right predication in an open list of novel cases – not necessarily all possible cases -- and thereby to give people something to talk and think about in those cases. Meaning sceptics assert that no finite set of examples determines the right application in any new case. Non-sceptics deny this, committing themselves to the claim that a relevant rule of correctness can be determined by how the subject responds to such a set of examples. But non-sceptics can think that finite sets of examples determine how a term should be applied in new cases without thinking that they can determine how it is to be applied in all conceivable new cases.

Third caveat. The response-dependent story of semantic and conceptual mastery does not entail that in using the term or concept learned people will be predicating, not the property at the source of the response – the property that unifies those instances towards which people are disposed to extrapolate from the original examples – but rather the relational property of typically arousing that response in human beings. That the predicate 'regular' is learned response-dependently does not mean that those who learn it think of the property that unifies regular shapes as one of looking a certain way to them: they may think of it as that property, the intrinsic property exemplified in such and such shapes. Of course some response-dependently mastered terms may pick out relational, in particular anthropocentric properties: the predicate 'funny', as used of jokes, surely picks out such a property. But the point is that the response-dependent story of semantic competence does not necessarily make basic terms of discourse anthropocentric in the same way (Jackson and Pettit 2002; cf. Peacocke 1993). Those who use a response-dependent term for certain items may do so in virtue of having the reference-fixing response and yet not be aware of that response; thus they can hardly be held to use the term in order to predicate the property of giving rise to that response.

Steps 6 and 7

The last two steps of the argument finally take us to the doctrine for which I want to use the name 'perspectivism'. Human discourse is not a single seamless web in which the basic terms are introduced and applied on the basis of a certain set of universally available responses, and other terms then defined by reference to those basic terms. Human discourse, rather, is a patchwork. The patchiness comes of the fact that the responses that facilitate the mastery of basic terms are driven by different, clustered sets of factors. As these factors cluster and divide from one another, so too do corresponding modes of discourse. And what perspectivism claims is precisely that in this sense there are a number of different discourses.

It is fairly obvious that human thought is structured around vocabularies, however loosely demarcated, such that we learn to master each in a manner that does not require an ability to define its terms in other vocabularies. We learn to master the language of belief and desire, and to use this language in explanation of one another's doings, without relying on definitions of its terms in the language of neuroscience – and vice versa. We learn to master the language of colour and beauty, sound and music, without relying on definitions of its terms in the language of physics – and vice versa. We learn to master the language of normative assessment, identifying justice and kindness and rightness and so on, without relying on definitions of its terms in the language of non-evaluative fact – and vice versa.

What the claim in the last two stages of the argument does is help to make sense of such palpable facts. Given the response-dependent ways in which we gain semantic mastery of basic terms, it is clearly possible that the factors tuning responses cluster in such a way that they give rise to relatively independent bodies of semantic competence and discourse. The sorts of examples given are good evidence that that possibility is realised. It is a matter of common observation that different discourses enjoy a high level of definitional autonomy in relation to one another and the response-dependent story of semantic mastery makes very good sense of that fact.

We need not consider here the question of how widespread discursive cleavages go, though we shall be suggesting further examples later. Some will think that many of the most interesting divergences occur only across different cultures, for example, while others will restrict themselves to divergences – such as those given in our illustrations – that we are as likely to find within any single culture. But I abstract from such undoubtedly interesting questions here.

One last comment. Different discourses are not only relatively independent in the sense explained; sometimes they are also mutually inhibiting. The best example is provided by the way that the use of ordinary intentional vocabulary to make sense of a psychological subject inhibits the use of neuroscientific vocabulary to do so; and vice versa. This example is nicely dramatised by Daniel Dennett (1979) when he asks us to imagine playing chess with a computer (cf. Davidson 1980; 1984). It is clear that if we are to play chess then we had better treat the computer in the vocabulary of intentional states: the language of beliefs, desires, intentions, and so on. It is clear that this is to do something different – to take a different stance, as Dennett says – from what we might do in looking at the electronic configuration of the computer and at how it is electronically configured to respond. We cannot simultaneously think of the computer in the two vocabularies: we have a choice between seeing it as an intentional subject, and playing chess, or seeing it as an electronic system and pursuing a purely predictive project.

I do not suppose that wherever we find relatively independent discourses, in the sense explained, we find discourses that are also mutually inhibiting. But it is a possibility worth keeping in mind. We shall see that it is often realised, and I suspect that many of those who think of themselves as perspectivists give it great prominence. Perspectives are mutually inhibiting or exclusive to the extent that you cannot see things from two angles at once, and so the very metaphor at the origin of perspectivist claims suggests the thesis of mutual inhibition as well as mutual independence.

Responses to perspectivism

How should we react to the fact that there are many different perspectives that we can assume, as we put ourselves under the control of different sorts of tuning factors? Sensitised to the effects of things on the senses, for example, and to the ways in which they command our emotions and desires, we can see a world of colour and beauty and value. Sensitised to the ways in which they fit and interact with other things – our bodies included – we can see a world of shape and volume and mass and causation. What does it say about the nature of the world that it can be depicted – and, we presume, accurately depicted – in such different ways?

One salient answer would say that the reason why representations are multiple is that reality itself is multiple. Corresponding to different forms of representation, so the picture would go, there are different realms of reality that relate in the manner of Cartesian mind and Cartesian matter, or the Platonic realm of sense and the Platonic realm of intellect, or perhaps Nelson Goodman's (1978) many worlds.

But the multiple-reality story will not attract many contemporary adherents. Some will refuse to endorse it on the grounds that it would make sense only if we had access to an external, god's-eye depiction of the things represented in different perspectives (Price 1997). Most will reject it, as I reject it, on the grounds that it fails to explain the intuitive fact that properties in the supposedly different realms are often superveniently connected: aesthetic properties supervene on perceptual, for example, evaluative properties supervene on descriptive and, according to most accounts, mental properties supervene on physical. There is no variation possible in supervenient properties without variation in the subvenient; and this sort of connection, however it is further articulated, is hard to square with a hypothesis of multiple, distinct realms.

A second salient answer would say that reality can be multiply represented without itself being multiple in the envisaged sense, because representation is constructive in character. Spelt out in cruder terms than adherents may favour, the idea is that reality is relatively indeterminate and that our more determinate representations vary as we choose to construct an image of it, now with this purpose in mind, now with that. Each image will be too determinate to be a faithful reflection of what there is and the different images can be seen as different constructions that we impose, for different pragmatic reasons, on the relatively unconstraining data that we confront in experience and experiment. There may not be a multiple reality, then, but there is multiple vision. There are many ways of seeing things – ways of seeing things that may not be easily reconciled (Nagel 1986: 88) – but there are not many ways things are.

This second alternative is most readily associated with Nietzsche's so-called perspectivism, at least on the received reading of that doctrine under which it posits an indeterminate nature (Leiter 1994). The alternative can be given different glosses, too, some idealist in character, some Kantian, some pragmatist, but no matter how it is presented, the approach is too mysterious for my own tastes. It offends against a natural sense of how representations ought to track reality, even if it does not violate intuitions in the outrageous manner of the multiple-reality account.

The third salient answer to the question of how the world can allow of multiple representations is reductivist in character. It refuses to multiply either reality or vision, arguing that some representations are more fundamental than others and that it must be possible, in principle, to identify a single form of representation that is the most basic of all those available. The reductivist idea is that we can think of reality as answering to this encompassing scheme of representation and make sense, at the same time, of how other, more specialised forms of representation are also possible.

I favour this third, reductivist answer to the question of how the world can allow itself to be multiply represented – but not in the stock versions. In stock versions of reductivism, the claims that we make within a special form of representation are analysed in such a way, so it is alleged, that we can see in principle how to construct equivalent and equipotent claims – if you like, translations – within a more comprehensive scheme; that, or they are explained away as the product of intelligible error. I do not think that reduction is always either translational, in any straightforward sense, or eliminative, as indeed many contemporary philosophers acknowledge (see Papineau 1995; Jackson 1998). I believe that once we recognise the multiple, perspectival nature of representation, we can identify a third variety of reduction that is neither translational nor eliminative. I shall call this third variety of reduction 'integrative'; integration is not as good as translation but it is a lot better, of course, than elimination.

The metaphor of visual perspective already points us towards the possibility of integrative reduction. It is possible for one visual perspective to be more encompassing than another in the sense that everything visible from the encompassed perspective will be visible from the encompassing but not vice versa; think, for example, of the balcony view of the front of the theatre and of the way in which this encompasses the view from the stalls. The encompassing perspective in such a case will deliver to us everything that is delivered in the encompassed, though not at the same angle and not perhaps at the same level of resolution; the ontology of the narrower scene will be included in the ontology of the wider. Consistently with this ontic reduction, the narrower perspective may retain a certain visual autonomy. It may be that you have to occupy that perspective in order to experience, or even vividly imagine, how things present themselves there. Although it is ontically reduced, as I shall say, the perspective remains heuristically indispensable; it is associated with a way of knowing how things are at the front of the theatre that is not available from the balcony.

By analogy, the idea in integrative reduction is that it may be possible to reduce one perspective to another in an ontic sense while allowing a certain autonomy – a heuristic autonomy – to the reduced perspective. The sort of representation available at the reduced perspective may not be capable of being replicated at the reducing; it may incorporate a way of knowing things that is lost at the reducing. The ontic reduction may be effected without a heuristic reduction.[1]

So far as there is ontic reduction, what is offered in the exercise envisaged is something more than elimination. But so far as it is not a form of heuristic reduction – so far as it loses the way of knowing associated with the reduced discourse – what is offered is something less than translation. The eliminative exercise involves neither ontic nor heuristic reduction; the translational exercise involves both. Integrative reduction is a half-way house, offering reduction of an ontic but not of a heuristic kind.

A model for the integrative reduction of perspectives

Consider the different representational perspectives associated with a purely non-indexical vocabulary, on the one hand, and a vocabulary that includes indexical terms on the other. Indexical terms identify times by reference to time of utterance as 'then' and 'now'; identify places by reference to place of utterance as 'there' and 'here'; identify people by reference to the identity of the utterer as 'I' and 'you'; and so on. Non-indexical terms can identify those very same times and places and persons but will do so without relying on the utterance in which they figure; they will pick them out by names, co-ordinates, calendars, clocks, and the like.

There is an obvious sense in which indexical facts – the facts which can be truly reported in sentences that use some indexical terms – reduce ontically to non-indexical; and yet in an equally obvious sense indexical facts remain heuristically irreducible. There is a clear sense, that is to say, in which the indexical perspective can be integrated with the non-indexical in a manner that parallels the way in which the view from the stalls can be integrated with the view from the balcony.

The ontic reducibility of indexical to non-indexical facts can be expressed as follows. Think of all the non-indexical facts that actually obtain, including facts about which speakers are at which venues at which times. Imagine now that we exactly replicate the actual world in these non-indexical respects, adding nothing on the way (Jackson 1998). Will the replicated world display all and only the indexical facts that obtain at the actual world? Of course it will. Let the non-indexical facts be carried over to the replica, and the indexical will travel at no extra expense. They will travel for free, because they depend on the non-indexical.

If there is any doubt about this, a little reflection will put it to rest. In replicating the non-indexical facts of the actual world, we replicate the positions of speakers in space and time. This is so for any indexical sentence that can be uttered by such a speaker; we ensure that things are such that it has the same truth-condition[2] and the same truth-value as the corresponding sentence in the actual world. But that is just to say that in replicating the non-indexical facts of the actual world, we replicate the indexical.

Here is another way of emphasising the ontic reducibility of indexical to non-indexical facts. Suppose that we are given complete non-indexical information on the nature of the world that we inhabit. Will that information leave open further possibilities of an indexical kind? Will it be consistent with the truth of any of a number of inconsistent indexical sentences? And will we have to wait, then, on information as to which of those sentences is true before the open possibilities are closed? Of course not. There can be no bare, indexical difference between possibilities. There are no indexical ways things can be that are fixed independently of how things are in non-indexical respects. No indexical difference without a non-indexical difference (Lewis 1990: 505).

If indexical facts are reducible in this ontic sense to non-indexical, it should be equally clear that they are not reducible in a heuristic sense. There is no way of interpreting indexical reports in non-indexical terms such that the mode of knowledge associated with the indexical reports can be equally well provided – even putting aside problems of extra complexity – by their non-indexical counterparts. The indexical perspective enjoys heuristic autonomy in relation to the non-indexical.

Consider any indexical report such as 'I am in Canberra' or 'It is 3.30 p.m. now' or 'Here is a tennis partner'. With such a sentence we can always specify when those words will express a truth. 'I am in Canberra' is true for any speaker S and any time, t, if and only if S is in Canberra at t. 'It is 3.30 p.m. now' is true for any speaker S and any time t if and only if it is 3.30 p.m. at the time of utterance. But while we can produce such biconditionals – and, as it happens, such a priori true biconditionals – we cannot claim to be able to use them in order to offer the speakers non-indexical ways of registering the things that are known, in the mode in which they are known, at the indexical perspective.

This point is made salient by John Perry's well-known argument that indexical modes of thinking are essential for agents (Lewis 1983, Essay 10; Perry 1979).[3] Take the indexically recognised facts that prompt me into action: this, in the way that my recognising that I am in Canberra prompts me to check my email or that my realising that it is 3.30 p.m. now prompts me to go to afternoon tea. Suppose that I could only recognise such facts in a wholly non-indexical mode: I could register that P. P. is in Canberra but not that I am in Canberra, and not that I am P.P.; I could register that afternoon tea is at 3.30 p.m. but not that it is now 3.30 p.m., and not that afternoon tea is now. In such a case, and regardless of the strength of my desires, my beliefs would be incapable of prompting me to action. For why should a belief about P.P.'s whereabouts lead me to do anything, short of knowing that I am P.P.? And why should a belief about what happens at 3.30 p.m. prompt me to take any initiative unless I recognise that it is now 3.30 p.m.?

The lesson is that while indexical facts reduce ontically to non-indexical ones, they are not heuristically reducible. In particular, they are not interpretable without serious loss in non-indexical language. They are not interpretable, indeed, without losing the very possibility of thinking in the manner of an intentional agent. The facts that are registered from the perspective of indexical language introduce nothing that is not encompassed by facts that are registered in the perspective of non-indexical: it is not as if indexical facts are something over and above non-indexical. Yet the indexical way of conceptualising and knowing those facts is not replaceable without serious loss by any non-indexical mode of conceptualisation or knowledge. Represent indexical facts in non-indexical fashion and they lose their normal profile and potency; they become incapable of heuristically mediating people's interventions as agents in the world around them.

In the non-indexical account of indexical phenomena, we are enabled to see from a non-indexical perspective how indexical sentences come to express truths. We can see in non-indexical terms that if the world is thus and so, then such and such indexical assertions will be true, such and such false. That is the sense in which the account offers a case of ontic reduction. But we are forced to admit at the same time that we cannot properly express or formulate those truths in non-indexical language: the indexical mode in which those truths are presented and known is not available at a non-indexical perspective. That is the sense in which the example illustrates heuristic irreducibility.

While we can offer a non-indexical reduction of what is sayable in indexical language, then, the reduction is neither translational nor eliminative. It is not eliminative, because it saves indexical truths. It is not translational, because it does not make available sentences that can serve in place of indexical sentences.

There is very little mystery, happily, about how this can be. The indexical perspective relies on a characteristic perspectival factor for the presentation of the propositions expressed in indexical sentences: this factor is the identity or whereabouts of the utterer, taken as such. Those propositions are presented either in terms that presuppose the identity of the speaker, as in talk of mine and yours and ours, or in terms that fix the speaker's location in space and time, as in talk of here and there, now and then. It is the identity of the actual speaker that fixes the reference of 'I' and 'you' and 'we', the location of the actual time and place of utterance that fixes the reference of 'now' and 'here'.

We can see from the non-indexical perspective how that factor plays a background role in the indexical presentation of potential belief contents; we see that for any speaker, for example, a sentence using 'I' will be true just in case it is true of that speaker. But we cannot put the factor into play in the same backgrounded way ourselves: we do not enjoy the appropriate speaker identity or location and, even if we did, the restriction to a non-indexical standpoint would prevent us from exploiting it. We can mention or identify the factor that serves a crucial role in indexical representation and we can explain how it serves in that role; but we cannot put the factor to any use ourselves (cf. Papineau 1995: 263).

The reduction of the indexical to the non-indexical serves as a model of integrative reduction, and as a model of how in principle perspectivism can be vindicated without recourse to the hypothesis of multiple vision or a multiple reality. What it illustrates is, broadly described, the following structure. There is a relatively narrow perspective, on the one side – the indexical perspective – at which certain patterns are salient: those expressed in the use of indexical sentences and those that must be grasped, by the Perry argument, for anyone intent on action. There is a relatively broad perspective, on the other – the non-indexical one – from which we can see how the narrow perspective operates and can recognise that it is directed to bona fide realities. But while the broader perspective enables us to recognise that the narrower is ontically reliable in this way, it does not itself enable us to represent directly the patterns – those expressed in the indexical sentences – that are salient at the narrower perspective. It enables us to see that there are such patterns to be seen at the narrower perspective but it does not make those patterns directly visible to someone restricted to the broad: to someone stuck with using non-indexical language. What such a person can represent is not the local pattern registered by 'It is now 3.30 p.m'. but rather the global pattern whereby a person using that sentence at a certain location in space-time will be detecting a corresponding local pattern.

The possibility of extending the model

My wager is that wherever we find the discursive divergences that perspectivism postulates, we will be able to reduce one discourse to another, or both discourses to a third, on the integrative pattern that the indexical case illustrates. I cannot vindicate my hypothesis definitively, of course, but I can do something to make it plausible. What I propose to do here is to show, in brisk outline, that integrative reducibility appears to be available – and to be all that is available – as between discourses that differ in a number of striking ways. I argue that figurative discourse is integratively reducible to non-figurative; practical representation, as I call it, to non-practical; and affective representation to non-affective.

Figurative and non-figurative representation

Some uses of language are meant only to amuse or shock or play but many, even many that are intended to elicit such effects, are representational in character: they are meant to convey a way things are. Among such linguistic representations of the world – among such conceptualisations, as I shall say – some are more or less figurative, some more or less non-figurative.

By non-figurative conceptualisation I mean that sort of representation in words or in concepts that relies solely on the pre-established meanings of the words used or, in the case of novel terms, on meanings that it explicitly introduces. Each referring expression picks out a determinate particular, each predicative expression a well-known property or a compound of well-known properties, each functor a familiar function, and so on, and each representational essay serves to convey something about the distribution of such properties across the universe of such particulars.

Conceptualisation is figurative precisely when the relation between the content represented – the story told, the scenario conveyed – and the meanings of the representational elements is not straightforward in this way. The words used are employed in such a way that the hearer or reader knows they are not meant literally – they are not meant in their established or stipulative senses – and knows that the content represented has to be identified on some basis other than by mere reference to those meanings.

Figurative representation may involve understatement, overstatement, insinuation, double entendre, irony, flattery or any of the many devices mapped in studies of rhetoric. But perhaps the best-known figurative form is provided by the example of metaphorical language. Here the user of the words says something so obviously false or banal that the hearer has to think twice and, relying on something more than the literal meanings of the words, has to identify a non-literal content. W.H. Auden writes in memory of Yeats: 'O all the instruments agree/The day of his death was a dark cold day'; but we do not hear in these words a banal report about the weather on a January day in Dublin, 1939. We make of the words – we hear in the words – a figurative message of a much more gripping and suggestive kind.[4]

We need not concern ourselves here with the detail of how such figurative speech is to be analysed. Without going into such detail, we can agree that there is a contrast between how it achieves representational ends and how non-figurative speech does so. Where non-figurative speech relies only on the pre-established or stipulative meanings of the words it employs, figurative – in particular, metaphorical – speech has to rely on something else besides. It puts into use, not just the literal meanings of the words, but the shared experience of the things and properties to which the words literally direct us and the shared sense of how those things and properties model items in the situation that the speaker is addressing. The dark and cold ascribed to the day of Yeats's death come to model features of the loss which his death entails. The instruments that chart the dark and cold come to stand for the indices by which we might measure that loss. And so on.

Given this dependence on experience, association and analogy we can see how the relation between figurative and non-figurative representation can fit the pattern distinctive of relations between narrower and broader perspectives. Figurative representation puts into play elements that non-figurative accounting cannot exploit in the same way, even though it can register and itemise those elements in the finest detail.

Experience, association, and analogy bring figurative speech to life for those who use and hear it; they let the language sing. There need be no mystery there from the literal perspective, since in principle the experience, association and analogy can be fully countenanced and their effects understood. But this ontic reducibility of the figurative to a literal perspective cannot be matched by anything deserving of the name of heuristic reduction. For the figurative mode of representation and knowledge requires experience, association and analogy to be put to use, not just to be given suitable mention in the ontological annals.

The point at issue here is familiar from discussions of analysis and art and the different methodologies of the two. The language of analysis, and the language of science in general, strives for exactitude and replicability and naturally embraces non-figurative modes of expression. The language of art strives for expression that is evocative rather than exact, and oriented to particularity rather than replicability. There is no end to what can be said in the language of analysis and there is no difficulty about the prospect of that language mapping all of the elements that go to make a language of art. Even if it can track the language of art in this way, still the language of analysis can never replace or exhaust it. The language of art may be worthy of celebration, indeed, for this very resistance to interpretation: for its irreplaceability. As Auden says in his lament for Yeats: 'poetry survives/In the valley of its saying where executives/Would never want to tamper'.

We argued before that, confined to a non-indexical perspective, we might be able to see exactly how indexical sentences could get to be true but we would not be able to express those truths in the same way: we would not be able to think about things in the manner required for intentional agency. We see now that the pattern is repeated in the case of figurative and non-figurative conceptualisation. Figurative representation is ontically unmysterious from the point of view of non-figurative but still it remains heuristically opaque. The discussion following will show that this pattern is repeated again and again.

Practical and non-practical representation

Imagine someone who argues regularly on these lines:

1 If every snake is timid, then that can't be a snake: look at the way it is tracking us.

2 That may have been a snake, for it certainly moved away quickly; but then again it may not, for lizards are timid as well.

3 If all snakes are timid, and this is a snake, then let's see what it does when I walk nearby: it ought to slither away into the bush.

4 You think that not all snakes are timid? OK, then show me one that is not.

5 Let's see whether there is any snake that will not be scared at the slightest sign of a human.

Imagine not just that the person argues regularly in this way about snakes, given their belief that all snakes are timid. Suppose that the person always argues on parallel lines for any belief of such a general or conditional kind. For arbitrary belief of the form 'All As are B', or 'If anything is an A, then it is also a B', they are prepared to argue from 'A' to 'B', from 'not B' to 'not A', from 'not A' to 'possibly B, possibly not B'; and from 'B' to 'possibly A, possibly not A'.

Does a person of this kind believe the principle associated in logic with ***modus ponens*** and ***modus tollens***? They do not believe it, we suppose, in the sense of being disposed in the manner of logicians to assent understandingly and sincerely to a sentence that formulates the principle; it may even be that they lack the words required to express it. But do they believe it in some other sense? Someone may say that they believe it in the sense of acting as if it were true. But that says little or nothing, for the principle is a necessary truth and it is hard to see how anyone could fail to act as if it were true (Stalnaker 1984). The question is whether the person believes the principle in a substantive sense akin to, but distinct from, the belief that goes with sincere assent.

We must surely agree that the person does believe the principle in such a sense. The person is a reasoning subject who invokes certain considerations as grounds for drawing certain conclusions: that is, for believing the conclusions in the assenting or judgemental manner. They endorse precisely the considerations that the principle makes it right for them to endorse: or at least that is what they do when they are being clear-headed and critical. They do not believe the principle judgementally, in the sense of being disposed to assent with understanding and sincerity to a sentence that formulates it; they believe it, as I shall say, in a practice-based or practical manner (Pettit 1998a).

Three features mark this mode of belief. First, the person is capable of recognising and understanding instances of reasoning that rely on the principle: for example, 'If every snake is timid, then that can't be a snake: look at the way it is tracking us'. Second, the person believes of precisely such instances that they are valid; even if he or she does not say that a suitable piece of reasoning is valid or that the conclusion follows or anything of that kind, they treat it as valid: they are disposed to produce the premises as reasons for maintaining the conclusion. Third, having such a universal and activated disposition to believe of instances of the principle that they are valid, the person counts as believing the principle ***in sensu diviso***; they manifest a case-by-case counterpart to the belief in the universal principle that all such instances are valid.

Once we see that it is possible for someone to believe and therefore to represent things in this practical manner, then it should be clear that many of the things we believe are represented for us in precisely this way. For our judgements are conducted all of the time on the basis of principles of reasoning that we rarely stop to spell out. Not just deductive principles of reasoning from judgement to judgement, as in the case of the ***modus ponens*** principle. Also the probabilistic principles that guide us in moving inductively from observation to observation, as we build up our picture of how things are in this or that situation. And also the principles that guide us in moving from perception to judgement and from judgement to the expectation of perception: the principle that takes us, for example, from the sensation of redness and the assumption that nothing is amiss to the judgement that something is red and from the judgement that something is red to the expectation that it will look red if nothing is amiss.

Nor is it just an accident that many of the things we believe are represented to us in a practical manner. Lewis Carroll's story of Achilles and the Tortoise shows that this is an inescapable feature of our beliefs. Achilles believes he can force the Tortoise to admit Z, given that he admits two premises, A and B, which logically imply Z. But the Tortoise insists that he needs more than A and B as premises; he also needs C: if A and B are true, Z must be true. Very well, says Achilles: you admit A and B and C, so now you must admit Z. No, replies the Tortoise, for I must also endorse another premise, D: if A and B and C are true, Z must be true. The lesson is that in order to reason from certain judgementally accepted premises to judgementally accepted conclusion, we have to believe in the practical way that such premises support such a conclusion; or if we take that as itself a premise, then we have to believe in the practical way that the new premises support the conclusion; and so on.

With our distinction in hand, it should be clear that in conceptually representing things to ourselves we may occupy either a relatively practical perspective or a relatively non-practical perspective. This distinction offers a parallel to the distinctions between indexical and non-indexical, figurative and non-figurative. If there is any contrast with those earlier dichotomies, it is simply that there is no such thing as the final, completely non-practical representation of anything: that is the lesson of Lewis Carroll's tale.

As the counterpart of those other distinctions, it should be no surprise to find that the divide between practical and non-practical conceptualisation holds out the same prospect of ontic but not heuristic reduction. It may be possible from a less practical point of view to take account of everything that shapes the practical representation of certain facts. But it is clearly not going to be possible to represent them and know them in the same way: some of the factors that the non-practical representation takes into account will be put to use in the practical representation in a way that cannot be replicated – though it may be simulated in imagination – at the less practical perspective.

Take, for example, the practical way in which most of us believe that various things follow from others. We believe them just so far as registering the entailing matters is not an inert affair; it involves seeing them as necessitating the entailed. The field of representation, as we might put it, is highly charged. These considerations, these matters of perception, just are reasons for making this or that judgement. Their status as reasons reaches out and grips our minds under the practical way of believing in the principles that countenance them as reasons. Thus the happy look on someone's face announces pleasure; the manifest weight of a load makes equally salient the strain that someone carrying it must feel; the clink of the billiard balls reveals immediately the causality at work in shifting the second, previously stationary ball; and so on. The world as practically represented is inferentially alive.

With any practical representation of things, it is clearly going to be possible to spell out the principles involved and to give a judgemental, less practical account of how things are depicted there. In this sense it is going to be possible to give an ontic reduction of the practical perspective to a perspective that is less practical. But equally clearly, it is going to be impossible at the less practical perspective to replicate the way things are seen at the practical one; it is going to be impossible to provide a heuristic reduction. For while we may be able at the less practical perspective to itemise the ratiocinative dispositions that play a role in the practical representation of things, we will not be able to put those dispositions to work in the same way. We will not be able to let the represented world come alive in the inferential manner associated with the presence of such dispositions. Someone might be quite willing to endorse in a non-practical way a variety of principles of reasoning but suffer from such an ailment that they are incapable of having those principles inform the way they see things; they might have to conduct their reasoning in a painful, step-by-step fashion.

Affective and non-affective representation

We mentioned in illustration of practical belief, that a person may believe in a principle of reasoning through being disposed to treat certain judgementally or perceptually registered contents in a suitable way: as reasons for drawing appropriate conclusions. But it is not only our judgements and perceptions that give us contents: that give us ways things seem. Our moods and emotions and desires may also give us corresponding appearances. Our gloom may make things seem hopeless, our embarrassment may make a situation excruciating, our pleasure may make a painting look beautiful, our desires may give various prospects an attractive or repulsive aspect.

Once we see this, we must admit the possibility that as people may argue from something's looking red in a situation where nothing is taken to be amiss to its actually being red, so they may argue in corresponding ways from parallel appearances. They may argue from something's looking attractive in a situation where nothing is taken to be amiss, for example, to its being the sort of thing that is genuinely attractive, even perhaps desirable (see Smith 1994). This observation shows that the category of practical belief may be broader than our earlier discussion suggested.

But it also shows something else. Suppose that I ordinarily come to judge that one or another option is desirable by reliance on the principle that links what looks attractive in situations where nothing is amiss to what is desirable. If I believe that principle in practical mode, then that means that I will ordinarily come to judge that one or another option is desirable through finding it attractive – through coming to desire it – in circumstances where nothing seems to be amiss. The attraction or non-attraction of options, in the absence of a belief that things are in any way amiss, will serve as the ordinary inferential prompt for judging that something is desirable or not desirable. My evaluative judgements will generally have an affective character. They will not leave me cold. They will be non-inductively associated with the presence of corresponding desires (Jackson and Pettit 1995). There will be an internal connection between evaluation – the fact of making an evaluation – and feeling.

Our observation points us, then, to a further divide among conceptual representations, apart from those between the indexical and non-indexical, the figurative and non-figurative, and the practical and non-practical. Conceptual representations of how things are may be affective, involving the presence of mood or emotion or desire or something of that kind. Or they may be more or less non-affective.

This divide, as we would expect, conforms to the now familiar pattern. We may fully understand from a non-affective point of view how it is that those who enjoy certain affections in the presence of things may come to represent those things in various patterns. We may be able to offer an ontic reduction of the affective perspective to a non-affective one. But that is not to say that we will be able to provide a heuristic reduction of the affective perspective. For the factors – the affections – which play a background role in the affective perspective, helping to shape how things are represented there, can only be taken into account in a foreground way within the non-affective. They can be itemised but they cannot be put to work; they can be mentioned but they cannot be used. We all realise from time to time that were nothing amiss we would indeed find a certain option the attractive one and that we do not actually find it attractive, because things are amiss in a certain way: we are at a low ebb and are suffering a temporary value-blindness. Someone who was locked into that position would be able to take full ontic account of how things are registered in ordinary judgements of value but they would not be able to provide themselves with the same mode of representation and knowledge.

Philosophy in a perspectivist perspective

The discussion of the integrative reducibility of the indexical to the non-indexical, and this last sketch of how the same structure is found in other cases, should help to make plausible the central claim I want to make about perspectivism. This is that while perspectivism is sound – while there are many different discourses, each enjoying a relative independence from the others – this does not mean either that reality is multiple or that our vision of reality is multiple. The discussion should establish the robust prospect that wherever there are divergent, discursive perspectives, it will be possible to provide an integrative reduction of one to the other, or of both to a third.

Our discussion does more, however, than just support this abstract claim and this challenge to those who think that perspectivism forces us to postulate a multiple reality or multiple vision. It points us towards an important lesson for philosophy in general, though not a lesson that I can develop fully here. This is that many stock objections to philosophical, broadly reductive accounts of certain phenomena may be misplaced; they may suppose that the aim is to provide translational reduction rather than reduction of an integrative sort.

Consider the complaint, associated in particular with Kripke's Wittgenstein, that there is nothing naturalistic in which rule-following can apparently consist: nothing that could normatively guide us, and guide us in an indefinite range of cases, in the way that we are apparently guided when we follow rules (Kripke 1982). Or consider the complaint that it is necessarily impossible to provide a compatibilist, naturalistic account of free will, for no matter how much indeterminacy is left by the physical world, and no matter how sophisticated the mechanisms that I embody, still there will be nothing available to answer to my sense that I can choose to do or not to do such and such (Van Inwagen 1983). Or consider, finally, the familiar reflection that however detailed an account we offer of the cognitive architecture and processing of human beings, or indeed of their neurobiological constitution, we cannot ever hope to identify consciousness thereby: such an account will never make sense of there being something it is like to be human and to enjoy human experience; it will never take us beyond physical dynamics (Chalmers 1996: 121).

The motif that recurs in complaints like these is: keep adding wood and all you'll have is wood. Keep adding non-normative factors and all you'll have is non-normative; keep adding non-autonomous causal loops and all you'll have is non-autonomous; keep adding non-conscious scanning devices and all you'll have is non-conscious. But from the point of view developed here, this criticism does not look to be overwhelming.

It is possible to account ontically for indexical representation and indexical facts – the matters reported indexically – in non-indexical terms; and yet we shouldn't think that such accounting will offer us anything that has the heuristic presence of the indexical. And a similar message holds for attempts to make sense of the figurative in terms of the non-figurative, the practical in terms of the non-practical, and the affective in terms of the non-affective. Why then should we expect attempts to account for normativity, free will or consciousness to give us a heuristic sense of that phenomenon? We might as sensibly hope for a proof of God's existence that would bring us to our knees at the conclusion.

I have argued elsewhere that the experience of normativity which characterises rule-following involves a representation of rules – say, a representation of the properties that we seek to track in predication – that puts into play two distinct background factors (Pettit 1990; 1996). One is a disposition to extrapolate in a certain way from various sets of exemplars of the predicate and so a tendency to see those exemplars as instances of a more or less determinate pattern or property. The other is a disposition to authorise other people's (or one's other selves') predicative responses, side by side with one's own, and to identify the property as that which is guaranteed to correspond with one's response only when things are such as to facilitate convergence between authorised respondents: only when things, as we theorists might say, are normal or even ideal (Pettit 1998b). Such an account of rule-following countenances a narrower perspective – that of the subject who instantiates the two dispositions – than the explanatory perspective from which it is offered. Even if the account is successful, then, it should not be surprising that it does not make palpable for someone at the explanatory perspective the pattern – a pattern expressed in the language of 'ought' and 'right' – that is allegedly salient to the subject of the dispositions.

What is true of such an account of rule-following ought to hold also for certain accounts of free will: for example, for those that follow in the tradition of Peter Strawson's (1982) paper on 'Freedom and Resentment', but without embracing multiple vision (see Pettit and Smith 1996). These accounts would seek to make impersonal, ontic sense of what free will involves, while emphasising that the ontic matters in question present themselves very differently within an involved interpersonal perspective. The key to the approach is the idea that what is distinctive about our experience of free will is not the facts registered in that experience so much as the heuristic mode in which they are registered. That mode of registering explains the grip that the notion of free will has on our minds and hearts.

As this lesson bears on the analysis of normativity and free will, so it may also bear on the possibility of accounting for consciousness. Suppose we can ontically reduce all that is tracked in conscious experience: say, in the experience of certain things as coloured. It ought not to be an objection that someone who comes to see red things for the first time, even someone who had total knowledge of the ontically reductive facts, will learn a new fact (Jackson 1982, 1986). We will say that however new the mode of presentation, what the person registers is a fact that is no more additional to the established facts than are indexical facts additional to non-indexical. But won't the person at least learn the new fact that old fact that can be presented in a second mode (Chalmers 1996: 142)? No. The person will also have known that previously; what he or she will achieve in coming to see things as red is something akin to a practical skill (Lewis 1990): access to the second mode of presentation.

In each of these cases, we can imagine a successful reduction of the phenomenon in question that does not do the equivalent of bringing those of us putting forward the reduction to our knees: that does not give us a palpable, engaged sense of normativity or free will or consciousness. The test of the success of the reduction will be that when we try to simulate the position of someone at the relevant, narrower perspective, we find ourselves forced to acknowledge that from that perspective things would be precisely as we experience them in our everyday sense of normativity, free will or consciousness. But it will require simulation of the narrower perspective to be persuaded of this; the reduction in itself will not carry a power of phenomenological conviction.

The lessons in the cases just mentioned combine to provide hope that there may be a comprehensive, relatively neutral scheme available within which we can give an integrative reduction of different perspectives and discourses.[5] Take physicalism, for example, understanding it to be the doctrine that if we replicate the microphysical aspect of the actual world, adding nothing on the way, then we will replicate all other aspects of the actual world as well; the actual world does not involve any particulars, any properties or any laws that are not fully provided for by microphysical arrangements (Jackson 1998; Pettit 1993). If physicalism is true, then it ought to be possible in microphysical terms to account ontically for all the phenomena there are. That may seem outlandish if ontic accounts are expected to be heuristic accounts, for the language of microphysics is clearly lacking in endless respects. It is not going to be indexical or figurative or affective in character, for example; and it is not going to be as practical as many other forms of representation. But once we recognise that ontic accounting does not mean heuristic accounting, the suggestion begins to make sense.

More than that, indeed, the suggestion begins to assume the form of a really challenging programme. Think of how philosophers have worked in recent years to make sense of how a purely physical brain and organism could come to be minded in the fashion of an intentional and thinking system. If physicalism is correct – or if it represents a useful working hypothesis – then that sort of investigation exemplifies a more general programme. This would consist in the attempt to make ontic sense of a variety of phenomena – phenomena like rule-following, free will and consciousness – under the assumption that the only materials allowed are those that we can reasonably expect to populate the microphysical realm.

To conclude, an objection to physicalism suggests that the microphysical perspective is more encompassing than the various perspectives we assume in

tracking colours and norms, for example, free capacities and conscious experiences. But in suggesting this, doesn't it thereby insinuate that really the phenomena tracked in those other viewpoints are chimeras? They seem to be there so long as we stick to a suitably narrow perspective but as soon as we move to a more comprehensive standpoint, they prove as elusive as the end of the rainbow: we reach out and lay our hands on nothing.

I strongly resist this objection. If physics is capable of tracking all that there is, and if it can countenance the phenomena that other narrower modes of representation track, then those other forms of representation also track realities; they can be seen as physics by other means.[6] If that does not seem to be a sufficiently elevated vocation for such ways of representing things, then we must remember that they are modes of access to the world charted in physics by means that physics itself cannot replicate. More than that, indeed, we must remember that they are modes of access to the physical world by means that are of inherent importance for creatures that are living, minded and socialised in the fashion of human beings. The world that would be fully theorised in the final physics only becomes a world fit for humans – a *Lebenswelt* – when it is presented in less encompassing perspectives.[7]

Notes

1 Some will find the word 'epistemic' more natural than 'heuristic' in this and in other contexts in the paper. But I avoid 'epistemic' for the following reason. Those who reject the 'knowledge argument' (Jackson 1982, 1986) say that there is no extra fact known by someone who experiences colour that cannot in principle be known by the person who knows all the physical facts but sees only in black-and-white. And yet such philosophers will typically stress that the person who experiences colour has a very different way of knowing some of those facts. What is different about their way of knowing can hardly be described as epistemic, given rejection of the knowledge argument, but it is readily described as heuristic.

2 I abstract from the difficulty that on some accounts of transworld identity the sentences will not refer to the same individuals and so will not in that strict sense have the same truth-conditions.

3 Ruth Millikan has made me aware that I need not endorse the detail of the Perry story – I may take a line that she herself would find more congenial – and yet use the story for my present purposes.

4 The comments I make here are probably consistent with the well-known position on metaphor defended by Donald Davidson (1978). He maintains that strictly a metaphorical sentence only has a literal truth-condition and in that sense a literal content, but he insists that by startling a hearer the sentence can serve to convey all sorts of other thoughts and reflections. But if Davidson's views are inconsistent with what I remark here, then that is probably the worse for those views. After all, it is surely manifest that metaphor does convey the sort of thing that can serve as the content of a belief. One thing that makes that manifest is that we can easily conditionalise on metaphor and that when we do so, we do not conditionalise on the literal content. If the day of Yeats's death was a dark, cold day, we can say, then he must have been a great poet. It would be a bad joke, and poor taste, to deduce instead: he must have died in winter.

5 My thanks to David Armstrong for pointing out that such a scheme may or may not be physicalist in character and so the lessons do not combine to support physicalism directly.

6 I abstract from the strategy introduced by David Lewis (1996) in his treatment of knowledge. Under this strategy claims at a narrower perspective fade into insignificance as we go to a broader, but I think that it is applicable only in very special cases – see O'Leary-Hawthorn and Pettit (1995).

7 This is the text of the Parcells Lecture, delivered at the University of Connecticut, Storrs in February 1999. I was greatly helped by comments received at the time, and later, from members of the audience, as I was helped by comments received when the paper was first presented at a seminar in the Australian National University in May 1997. My thanks for written comments to Andrew Gleeson, Michael Esfeld, Barry Hindess, Jakob Hohwy, Alex Miller, and a number of anonymous reviewers, and for separate, very useful discussions of the material to Tim Crane, Frank Jackson, Ruth Millikan, Michael Smith and, in particular, Huw Price.

Part II

VALUE PLURALISM AND LIBERALISM