Chapter 6

Intentionality: Outsides

> If a lion could talk, we could not understand him.
> — Wittgenstein, *Philosophical Investigations*

> 'Ouch' is a one-word sentence which a man may volunteer from time to time by way of laconic comment on the passing show.
> — W.V.O. Quine, *Word and Object*

In the previous chapter I discussed the problem of carving up the insides of an agent, of picking out its beliefs and desires, in order to make sense of its behaviour. But in order to do this we also have to solve the symmetrical problem carving up the outsides, i.e. picking out the objects of its environment. We not only want to know what the agent is thinking, but also what kinds of things it is thinking *of*.

This problem is often ignored since the only philosophers who worry about relating things-in-the-head to things-outside — i.e. realists — also tend to be the ones who assume that there is a fixed list of Objects in the world, that science will in the end tell them what those Objects are, and if the contents of the agent's thoughts are not on that list then they are just plain wrong (or not thinking about anything real at all). Therefore the job of discovering what kinds of things an agent may be thinking of is a job for natural scientists, not psychologists or philosophers. On the other hand Wittgenstein and Rorty notice the problem, but assume there is no solution since the only way to know what kinds of things a lion is thinking of is to *be* a lion. In this chapter I argue that there is a substantive problem here *and* that it is soluble from a third-person perspective.

6.1 Sense and Reference

Consider the following two examples. The first is Putnam's twin-earth experiment (1975). Suppose Jean is transported in her sleep from earth to twin-earth. Twin-earth is *exactly* like the original except that water is made out of XYZ not H_2O. However the stuff still looks and tastes the same, and so as far as Jean is concerned (not being a chemist) there is no difference. Putnam introduced this example to prove how there is more to the contents of our beliefs than the role that they play in our heads: the properties of the stuff that Jean calls 'water' have changed even though Jean's

thoughts about it have not. The aspect of Jean's thoughts about water that remains the same on earth and twin-earth is the 'narrow' content, and the aspect that has changed is the 'broad' content.

The second is Frege's example of the terms 'Evening Star' and 'Morning Star' (1892). These terms meant different things to the ancients but later astronomers discovered that they both referred to the same thing, namely Venus. Thus in Putnam's example there are two distinct referents (XYZ and H_2O) which Jean grasps using a single term with a single sense (Water). Whilst in Frege's example there is a single referent (Venus) that was grasped using two terms with distinct senses (Evening Star and Morning Star). Thus Frege and Putnam are trying to draw the same distinction between those aspects of meaning that play a role in the head (sense and narrow content, respectively) and those that do not (reference and broad content)[1].

The problem that Frege and Putnam's examples generate is this: Jean did not notice the difference between water and twin-water, and the ancients did not see the link between the Evening Star and the Morning Star. We are only able to draw the distinction between sense and reference (or narrow and broad content) in these cases because we take a kind of God's-eye view of the situation from which we are aware of things that the ancients, and Jean, were not. But for all we know there may be Higher beings that are making up the same kind of thought experiments about us ('imagine a group of people who stupidly think that water is H_2O, when of course modern superphysics shows that this was just a crude approximation'). Therefore our talk about the reference of water 'being' H_2O is just as sense-laden as Jean's take on the world. If referents are supposed to be independent of us, then how can we talk about them? As Putnam argues, '*What objects does the world consist of?* is a question that it only makes sense to ask *within* a theory or description' (1981, p49). Or, as Nabokov put it, 'reality' is a term that means nothing except when in quotes.

Now it must be acknowledged that many realistically-minded philosophers cannot see this problem, or — and this amounts to the same thing — they do not think that it matters. Why should it be necessary to be able to talk about the referent of our thoughts in a way that does not use our vocabulary of thoughts? Why should we require that it is possible to 'reduce' the referent, or state it in other terms? Such philosophers would argue we are talking about the referent simply by mentioning it, and that the thing we are talking about is independent of our way of talking. But the problem starts to manifest itself as soon as we try to prefer one way of describing the world over another. Frege would conclude, for example, that the term Venus is better in some way than either Evening or Morning Star; and Putnam's thought experiment seems to suggest that Jean would have been more accurate in describing the wet stuff in her world(s) as H_2O and XYZ respectively. But why?

Frege did not have a solution to this problem but he did have a way round it, a strategy that subsequently became fundamental to most analytic philosophy. Frege starts from Kant, and in particular Kant's assumption about what the world is 'really' like; namely that it is divided into independent objects-in-themselves which possess various properties. Frege concludes that if a language is to be scientifically respectable then its structure must reflect this essential structure of the world (Dummett, 1991b, ch20). This means that (scientifically respectable) sentences must be expressible in predicate-subject form, $P(x)$, where x refers to an object defined independently of that sentence, and P a property that is predicated of it:

[1] Fodor (1994) also draws a comparison between these two examples, though for a slightly different purpose.

> Statements in general, just like equations or inequalities or expressions in Analysis, can be imagined to be split up into two parts; one complete in itself, and the other in need of supplementation, or 'unsaturated'. Thus, e.g., we split up the sentence
>
> 'Caesar conquered Gaul'
>
> into 'Caesar' and 'conquered Gaul'. The second part is 'unsaturated' — it contains an empty place; only when this place is filled up with a proper name, or with an expression that replaces a proper name, does a complete sense appear.
>
> ...When we have thus admitted objects without restriction as arguments and values of functions, the question arises what it is that we are here calling an object. ...Here I can only say briefly: An object is anything that is not a function, so that an expression for it does not contain any empty place. (1891)

Frege's condition — that truth-bearing sentences must involve predication over a 'saturated' object term with prior reference — recurs in many forms. We find it in Russell's claim that a subject cannot make a judgement about something unless they can know which object their judgement is about; i.e. that the subject can refer to that object independently of that particular predication (1905). We also find the same assumption in Evan's *generality constraint* (1982), Fodor and Pylyshyn's criterion of *systematicity* (1988), and Millikan's condition of *propositional structure* (1984). The assumption is the same in each case: a subject cannot predicate a property, *P*, of an object, *x*, unless they can equally well predicate any other properties, *Q* or *R*, of *x*, and predicate *P* of any other objects, *y* or *z*. The same assumption also lies behind the argument within South Coast AI (section 4.3) that information-bearing functional states only count as representations to the extent that they are part of a more general symbol system. When Brooks, Beer, Harvey and Wheeler declare that one can have 'intelligence without representations' they mean representations that meet the generality constraint.

How does Frege's (or Russell's, or Evans') condition help us avoid the problem of reference? Compare the two sentences

1. The referent of the term 'the Morning Star' is Venus.
2. The referent of the term 'Venus' is the Morning Star.

Now strictly speaking both sentences are meaningless since there is no non-circular way of talking about the referent of a term. Nonetheless Frege *et al* give us a reason for *preferring* the first sentence to the second. The reason is that 'Venus' is a more *objective* term than 'the Morning Star' because the term 'Venus' may be used independently of when in the sky it appears, whereas 'the Morning Star' is more closely tied to particular observation conditions. For example, the sentence 'the Morning Star appears in the morning' is (almost) tautologous, but the sentence 'Venus appears in the morning' is not. In short, the term 'Venus' has a better claim to be part of an ideal scientific language than 'the Morning Star'. It is this that justifies the assertion of 1 rather than 2. Of course our thoughts may never come into ultimate correspondence with the way that the world 'really' is (and we would not know it even if it did), but we can still tell when we are getting closer. It might turn out that we were mistaken about Venus being a 'real' object all along, but the discovery of facts about Venus that are independent of its position in the sky is evidence that we are not. It is interesting to note that even the most hardened realist-minded scientists rarely claim

that they actually know the truth, instead they usually only claim that we can get better and better approximations to it. And the closeness of that approximation is judged according to how well our thoughts meet Frege's conditions[2].

Frege avoids the problem of reference, but does not solve it. He gives powerful reasons for preferring one system of description to another but these reasons are still based on an unjustified Kantian assumption about the essential structure of the world. The fundamental objection remains that the only things that our minds have access to — i.e. the only things that play a role in our heads — are senses. Frege defines sense as the mode of presentation of a referent to a mind but there is an alternative tradition, starting with the later Wittgenstein and including Rorty and the later Putnam, that defines sense independently of any notion of reference. According to this tradition 'the sense of an expression consists in its role within the complex social practice constituting the communal use of the language ... An individual speaker's grasp of that sense then becomes one ingredient in his ability, acquired by training, to engage in that practice' (Dummett, 1991a, p17). According to this tradition Truth does not lie in a correspondence between things-in-themselves and things-in-the-head (or expressions in a language), but in the ability of an individual to use language successfully.

For example, suppose that Jean walks into a bar and asks for a glass of water, but when she gets it she complains that the glass does not just contain water but also traces of mineral salts, bubbles of carbon monoxide, and a slice of lemon. Is it a glass of water? Clearly yes, since the norms of correct practice for bars define what constitutes 'a glass of water' in that context; different norms apply in a chemistry lab, where it is not normal (i.e. correct) practice to add ice and a slice to beakers of water.[3] Thus we can discuss the Truth of the use of terms like 'water' without recourse to metaphysical debates about what water 'really' is[4].

Can we do without a notion of reference, as Wittgenstein *et al* maintain? The sole reason for clinging onto reference is that we need some way of explaining why some representational vocabularies are more successful for certain purposes than others. Why, for example, does passing an electric current through water generate hydrogen and oxygen? Unless we believe that water *really is* made of H_2O then we are lost for an explanation. It is only the possibility of realism that makes the success of science non-miraculous, as Putnam put it (1973). This point is the mirror-image of that raised in the previous chapter: if we want to use the concept of belief to explain behaviour then beliefs must be instantiated in the brain mechanism underlying that behaviour; similarly, if we want to use reference to explain the success of the behaviour that those beliefs play a role in, then the reference of those beliefs must be instantiated in the world outside the head. But is there any non-circular way of talking about the reference of our thoughts?

There are two steps to breaking out of this circle. The first step is to adopt a third-person perspective. It is impossible to directly observe the world of reference outside our own head[5], but we can observe the world that surrounds other people's. But then there is still the problem

[2] I believe that one consequence of this is that mathematical physics, amongst all the sciences and rival systems of thought, is seen as having the best chance of grasping the real structure of the world since it approaches most closely to the Fregean ideal. But this is another question.

[3] The origins of these criteria of success and correctness will be the subject of chapters 8–11.

[4] Winograd and Flores (1986) use another water-based example to make a similar point.

[5] Those realistically-minded philosophers who do not acknowledge the fundamental problem of reference that lies behind this chapter will probably deny this claim; but I have a feeling I will have lost them a long time ago anyway.

of picking out the referents of their thoughts. We will have to use some vocabulary, some set of concepts, to pick out the objects that they are thinking of, and are we not trapped in using our own? What kinds of things would a lion talk of, if it could? The solution is to define a new vocabulary for picking out the objects in other agents' worlds, and one that is not based on our own concepts. This is the job that Evans' concept of Non-Conceptual Content (NCC) can do.

6.2 Non-Conceptual Content

Non-conceptual content is content that is characterised using concepts that the agent having the thought does not necessarily possess. For example, suppose we hear a sound coming a certain position in space and we turn our head to see where the sound was coming from. Looking back on the experience we may describe the reference of our state of mind in Fregean, conceptual, objective terms as 'a sound source located at position *X*', but this thought probably never occurred to us at the time. We just turned to see what it was. Evans suggests that it is more accurate to describe the contents of our thoughts in terms of our *activity*:

> What is involved in a subject's hearing a sound as coming from such and such a position in space? ...When we hear a sound as coming from a certain direction, we do not have to think or calculate which way to turn our heads (say) in order to look for the source of the sound. If we did have to do so, then it ought to be possible for two people to hear a sound as coming from the same direction and yet to be disposed to do quite different things in reacting to the sound, because of differences in their calculations. Since this does not appear to make sense we must say that having spatially significant perceptual information consists at least partially in being disposed to do various things.

Thus the content of the thought is better described as something like 'a sound coming from a direction that would be foveated if we turned our head *so*'. This description of the content uses concepts ('direction', 'foveate', etc) that we are not assuming that the agent used at the time, or even possesses, and instead makes essential reference to their ability to act in the world ('turn our head *so*') — described by Cussins as 'the realm of embodiment'.

Evans originally developed the theory of NCC to better describe our first-person perceptual experience of the world, and this is largely how the theory has subsequently been used by Peacocke (1992), Cussins (1992a), Crane (1992) and others. But it can also be used from a third-person perspective to describe the thought processes of agents without making any assumptions of whether they are consciously aware of them or not (Bermúdez, 1995)(Chrisley, 1995). Consider the example of a frog striking at flies. What is the frog thinking of when it does this? Of course in one sense it is not thinking of anything at all. If it has any consciousness then it is a very limited one. On the ladder of cognitive complexity it is only one step up from my reluctant car. Nonetheless it uses internal states that bear information about the environment in order to co-ordinate its interactions with the world — i.e. very simple representations — and so we can usefully explain its activity by ascribing it with simple proto-beliefs[6]. But what are these proto-beliefs *of*? From our Gods-eye view of the frog we can see that the things that it is striking at are flies. An entomologist

[6]If you are not happy with this example then we could take a few more steps up the ladder. How about a lion? A chimp? A new-born baby? All we need is an example of an agent whose behaviour we can usefully explain using psychological talk but whose conscious experience, if it has one, is very different from our own.

could go one step further and identify the individual *species* of fly. But it is obvious that these fine distinctions in carving up the world play no role in the frog's activity. A frog does not have the concept 'fly', let alone a concept of particular species. As far as it is behaviourally concerned all it perceives are 'things that it should strike at like *so* and eat'. This is the content of its belief, described using concepts that we are not pre-supposing that the frog itself possesses.

Evans' analysis has similarities with Gibson's ecological theory of perception (1966)(1979)[7]. Gibson argues that the basic things we perceive are not objects as the reductionist imagines them, but *affordances*. An affordance, as defined by Gibson, is the way in which the environment can play a role in the behaviour of a creature. For example, the surface of a pond provides the affordance of 'something to walk on' for a pond skater, but not for a human. An affordance describes an environment *for*, or with respect to the behaviour of, a creature. It describes the world *of* the creature, and not the world as viewed from nowhere — or, rather, from the point of view of no body. An affordance is thus a relational property of an environment, defined with respect to a specific behaviour of a specific organism. The same object can provide many different affordances for a single creature and, conversely, different objects can offer the same affordance. Flies and bees both provide the affordance of eatability for a frog, while a pond provides both the affordance of spawning, and of escaping from terrestrial predators.

Specifying contents in terms of affordances turns the standard analytic account of perception on its head. This account was inherited from the British empiricists by the logical positivists and passed, *via* Carnap, to cognitivism, computationalism and East Coast AI, where it found canonical expression in the work of Marr on vision (1982). According to this account, perception starts when we form an internal map of the objects in our world from sense-data. These perceived objects are then categorised and attributed with properties so that we can plan our behaviour. Therefore the ability to act in the world is built upon a more basic ability to perceive objects. But according to Evans and Gibson perception is not primarily the ability to form an objective map of one's environment, but the ability to act within it[8]. Most of the time we are actively engaged in the world rather than considering it passively. (Philosophers are the exception to this rule, which explains why they are so fond of the standard model of perception.) Of course we are also able to perceive objects passively, to categorise them and predicate properties of them, but this passive ability to perceive objects is built on a more basic ability to perceive affordances in our environment.

For example, when we reach for a saucepan or avoid tripping over the cat we do not primarily perceive them as categorised and labelled objects but as something like 'things to cook with' or 'things to avoid'. Indeed, even these descriptions of the contents of our thoughts are misleading since, by definition, it is not possible to give an accurate translation of the content of a non-conceptual thought using plain English concepts. One alternative way of describing this content is to hyphenate the description (e.g. 'thing-to-cook-with'), to show that the content should be understood as a unified whole rather than a construction of sub-concepts. But the canonical way of describing such contents is to describe the activity they play a role in. Thus, for example, it is more accurate to say that our thoughts about the saucepan at the time was comprised of a correct recognition that the environment enabled us to boil potatoes by moving and acting in certain ways.

[7] And Rowlands (1997) shows how this analysis fits neatly into an evolutionary framework.

[8] Gibson's analysis of perception is often seen as incompatible with a referential theory of mind, but Sloman gives an example of how this is not necessarily the case (1989).

One consequence of this is that in order to communicate the content of a (non-conceptual) thought from one head to another it is not sufficient that they share the same language; they must also share the same body (or, at least, the same abilities to act in their environment). Someone who cannot cook, for example, cannot really know what a skilled chef means by 'a saucepan' because they cannot know what a saucepan means to them.

NCC enables us to describe the contents of thoughts using concepts that the agent having the thought does not necessarily possess. This may be because the agent does not have any concepts at all (as in the case of the frog), it may be because the agent has those concepts but did not use them at the time (as in the case of our perception of cooking implements), but it may also be because the agent uses concepts that are different from our own. Consider Quine's example of the linguist trying to understand the language of a native tribe (1960). The linguist observes that whenever a native sees a rabbit they point and say *gavagai*, but how should they use this evidence to decide what 'gavagai' means? Quine argues that the meaning of 'gavagai' is underdetermined by the observed behaviour. Just by observing the native pointing to rabbits the linguist does not have enough evidence to decide whether 'gavagai' means the same as 'rabbit', or if it means something like 'undisconnected rabbit parts' — you cannot have one without the other, and so there seems to be no principled way of choosing between the two possible meanings. But there is more to the meaning of 'gavagai' than its observation conditions. There is also the role that rabbits play in the life of the native speaker. And it is this aspect of meaning that we can use the technique of NCC to describe.

For example, suppose that the linguist discovered that the native had access to a genetics lab, and was fond of sequencing the DNA of any animals that she came across. This would be evidence that the native's concept of gavagai had a similar meaning to that of an English-speaking biologist's concept of rabbit — something along the lines of 'members of an inter-fertile species of rodents identified through their possession of a wild-type genome of *XYZ*'. On the other hand if the linguist found that rabbits were no more than a source of food and skin to the native, then the meaning of *gavagai* would change accordingly. Thus the linguist can try to understand the meaning of a term by understanding the role that it plays in the life of the native speaker, rather than by trying to find a concept in her own vocabulary that is an exact translation of it. Indeed it is unlikely that, in the former case, the English-speaking linguist would have a concept that is a synonym of 'gavagai' for the same reason that she could not quite grasp an English-speaking biologist's concept of 'rabbit'. The obstacle preventing an accurate translation is that the linguist and biologist have different ways of interacting with rabbits, rather than different vocabularies.

The meaning of a term is defined by the role it plays in an agent's overall activity — its being-in-the-world, form of life, existence, world, praxis, practical discourse, lifestyle, or mode of production, depending on one's choice of existential philosopher[9]. Words are tools, and their meaning is defined through how they are used. Now when we describe the use of a tool, we are not trying to 'translate' that tool into words. When we describe how a hammer is used we do not suppose that there is any hammering going on in our heads. Similarly, when the linguist describes the native biologist's concept of 'gavagai' as 'a member of a species . . . ' she is using her concepts

[9] In chapters 10 and 11 I argue why the latter term is more useful by discussing its role within the evolution of culture. But this is another question.

to describe the activity of the agent through which the meaning of 'gavagai' is defined for them. The linguist is *not* thereby assuming that the native shares the same concepts, any more than we assumed that the frog understood the concept of 'striking'. Thus NCC provides a way of carving up the world in order to make sense of an agent's behaviour that does not involve foisting the distinctions made by our concepts onto them.

6.3 Affordances and Objects

According to the Fregean tradition, conceptual thoughts refer to objects-in-themselves — these comprise what Cussins describes as the 'realm of reference'. This tradition then produced the problem of talking about the realm of reference independently of our own way of talking. NCC seems to offer a solution to this problem by talking about the content of thoughts in terms of physical activity in a way that does not pre-suppose that the holder of the thoughts has the same concepts as us. But what do non-conceptual thoughts refer to? (Or, what are the referents of mental states whose contents are described non-conceptually?) Here opinions differ. Cussins, following Evans, argues that describing the contents of thoughts non-conceptually is a way of *avoiding* defining their reference:

> *What Evans saw was how to pull apart the specification of content from the specification of reference or truth.* If a canonical specification of a content need not be a specification of a truth condition, then canonical specification of a content which refers to the realm of embodiment does not entail the evident falsehood that the truth of the content depends on the character of the realm of embodiment. ... For Evans, truth conditions are fixed by the realm of reference, and not by the realm of embodiment; but the cognitive significance of representation is fixed by the realm of embodiment, and not by the realm of reference. (1992a, p656, original emphasis)

This argument stems from the Fregean assumption that the 'real' world — i.e. the realm of reference — is 'really' divided up into independent objects-in-themselves. Therefore if the contents of thoughts fail to carve up the world in this way then they cannot be referring to that realm. There are still objects that non-conceptual thoughts will be true of, but these objects play no role in the life of the agent *per se* (they will not be 'cognitively significant'). The real world remains hidden from the agent trapped in its realm of embodiment. Thus the frog only 'knows' about 'things it should strike at', not the various species of fly that infest the real world of the realm of reference. Frege's problem of reference remains unsolved but Evans' achievement, as far as Cussins is concerned, is to show how we can have a theory of content despite this.

Epistemology recapitulates ontology. In other words our theory of how we know the world depends on our theory of what the world is like. In chapter 2 I tried to loosen our Kantian assumptions about the essential structure of the world, and this has implications for Fregean assumptions about reference. In particular I argued that the world is *not* made up of independent and prior classes of objects-in-themselves which then come together to form larger wholes. All things on all occasions exist in contexts, therefore the boundary between an object and its world is not an ontological given; and nor are the criteria that make two objects 'the same' (or of the same type)[10]. We

[10]This also implies that *sets* of objects are not ontologically given — this corollary was important for the discussion concerning the individuation of theoretical terms on page 52n.

do not carve the world at its joints, rather we define joints through acts of carving. Moreover we carve the world in particular ways because those ways yield objects that have properties that are useful for us; in other words because those objects provide affordances. The frog picks out 'eatable things' because doing so helps it survive. We pick out 'species of fly' or 'gas molecules' or 'washing machines' because doing so helps us understand evolution, the behaviour of bulk gasses, and getting clothes clean, respectively. An important criterion of the scientific way of carving the world is that it should pick out objects that have properties that are relatively constant across contexts; but this does not mean that the resulting objects are 'more real' than those carved out for other reasons. For example, the property of being 'an eatable thing for a frog' is just as 'real' and objective as the property of being 'a member of a species defined by a biologist'.

Nonetheless, as humans, it seems we *can* perceive objects independently of any particular affordance that they offer. Indeed for most philosophers this has seemed like the most basic kind of perception. Gibson argues that when we are cooking we directly perceive that a saucepan offers the affordances we need in order to make a sauce, and the thought 'that object is a saucepan' may not enter our heads. But it *may* enter our heads when we pause and look at our kitchen passively. We can identify the objects present, give them names, and think what predicates are true of them. Where does this ability to perceive objects *as* objects, independently of use, come from? And why does it seem so basic to perception?

A frog does not contemplate the black dot floating in front of it, decide that it may be something worth eating, and strike. All it perceives is an eatable thing *there*. Animals do not contemplate the world passively but are engaged with it in a constant pursuit of the four F's. But as animals become more complex their behavioural repertoire grows. They become more flexible, and use their environment in different ways. Humans are the most extreme product of this process. With our free hands, opposing thumbs, and bulging cortexes we can learn to work on our environment such that it offers affordances that support a virtually unlimited range of behaviours[11]. We have even evolved the ability to learn how to manipulate objects in our heads, rather than in our hands; to imagine and consciously explore the affordances that they offer beyond the immediately F-able, and to share our findings through language. As agents become more behaviourally sophisticated their perception of objects becomes less tied to any particular behaviour that they may play a role in. Thus the perceptual space of humans becomes more objectified and conceptual as the richness of our interactions with the world increases[12]

Frogs, for example, only eat flies; and the only thing they do with flies it eat them. Frogs cannot be said to predicate the property Food of the object-class Fly, since neither exists for a frog independently of the other. Frogs cannot have any concept of 'food' separate from that of 'fly'. Humans, on the other hand, are spectacularly omnivorous. There is scarcely a single biological entity that we cannot turn into food, therefore we can support a concept of 'food' independently of any particular food*stuff*. We can also do things with, say, potatoes other than eat them — such as paint them, carve them into potato heads, or sell them. Therefore we can support a concept of 'potato' independently of the predicate 'eatable thing'.

[11] The argument that the evolution of human brains and psychology is fundamentally grounded in our ability to manipulate the environment, rather than our ability to manipulate concepts, was central to Engels' pamphlet *The Part Played by Labour in the Transition from Ape to Man* (1987) — see also *Posture Maketh the Man* (Gould, 1978, ch26).

[12] For example Cussins discusses how our objective perception of Euclidean space evolves with our ability to navigate that space (1992a) — see also Bennett (1996).

Traditional societies tend to be fairly fixed in the ways that they interact with the world. Conventions, taboos, and tradition dictate what use each type of object may be put to. But the birth of capitalism swept away all fixed ways of interacting with the world. Everything was up for grabs — or rather everything could be grasped in any conceivable way. The only limit on our ability to interact with the world — and hence to perceive it — is now the limits of those objects themselves. As Marx, in the Communist Manifesto, put it:

> Conservation of the old modes of production in unaltered form was the first condition of existence of all earlier industrial classes. Constant revolutionising of production, uninterrupted disturbance of all social conditions, everlasting uncertainty and agitation distinguish the bourgeois epoch from all earlier ones. All fixed, fast-frozen relations, with their trains of ancient and venerable prejudices and opinions, are swept away, all new-formed ones become antiquated before they can ossify. All that is solid melts into air, all that is holy is profaned, and man is at last compelled to face, with sober senses, his real conditions of life and his relations with his kind.

Our ways of carving the world are not dictated by fixed ways of interacting with the world. But this does not mean those ways of carving are *independent* of our interaction with the world, as the picture theorists argued. The perception of objects — i.e. the identification of references — can never be completely independent of the activity that those objects play a role in because the act of identification itself is based on activity. As Dummett argues

> We have the notion of the bearer of a name, and the conception of a predicate's being true or false of an object, in advance of constructing a semantic account of our language in order to analyse its working, because these are embodied in quite primitive linguistic performances; our acquiring them is part of our learning to use our language. Both are born of the practice of ostension, that is, from our possession, in the use of a demonstrative accompanied by a pointing gesture, of another means than the employment of a name for picking out a concrete object. By means of a recognition statement (a statement of the form 'This is *a*'), we are accustomed to identifying an object as the bearer of a name; by means of ostensive predications (statements of the form 'This is *F*'), we are accustomed to applying predicates to objects picked out ostensively. To say that the referent of a name is its bearer, and that the referent is what we speak about, is in effect to say that the semantic roles of proper names and of simple predicates should be understood in relation to these fundamental practices: it is precisely because of our thorough familiarity with these basic linguistic practices that the notion of reference supplies us, as soon as it is introduced, with so definite and readily acceptable a picture of the semantic roles of at least the simplest logical types of expressions. (1981, p406)

6.4 Conclusion

In these two chapters I have argued that successful intentional behaviour depends on an accurate correspondence between things in the head (representations) and things outside (objects). In other words I am a realist and a representationalist. The pragmatist's objection to this kind of realist representationalism is that he thinks that it bases its explanations of success on a list of Objects as they Really Are. But, he argues, we have no way of determining what these Objects are beyond the pragmatic success of our own theories about them:

> There is no independent test of the accuracy of correspondence . . . The representationalist's attempt to explain the success of physics and the failure of astrology is bound to be merely an empty compliment unless we can attain what [Putnam] calls a God's eye standpoint — one which has somehow broken out of our language and our beliefs and tested them against something known without their aid. But we have no idea what it would be like to be at that standpoint. . . . My principal motive is the belief that we can still make admirable sense of our lives even if we cease to have what Nagel calls "an ambition of transcendence" (Rorty, 1991a, p6,12)

Of course we cannot have a Gods-eye view of ourselves, but the third-person perspective can give us a kind of God's eye view of the relationship between another agent and *their* world — as long as we solve two problems. The first is to identify the thoughts inside the agent's head. The solution to this, I suggested, requires that we *look* inside their heads: the correct intentional description of an agent depends not only on their external behaviour but also upon the mechanisms that underlie that behaviour. The second problem is to identify the objects of their world. And this, I suggested, requires that we do not pre-suppose our vocabulary — our way of carving up the world — will correctly carve up theirs. We may not be able to use *our* object-concepts to describe the referents of *their* thoughts. The alternative to using our concepts is to pick out those referents on the basis of their activity: their beliefs may still involve concepts, but they won't necessarily involve *our* concepts. Once we take both of these steps then we can reinstate truth — i.e. a relationship between things-in-the-head and things-out-there — as a notion capable of genuinely explaining the success of another agent's actions, even though we cannot use it directly to explain the success of our own. But if truth explains the success of the actions of the other agents we see around us, then it seems like a reasonable induction to suppose that it would also be capable of explaining the success of our own. (This is the third-person equivalent of the Other Minds Problem: instead of wondering whether other people have consciousness like us, we wonder whether we have the ability to represent an external world like them.) Rorty is correct that we cannot transcend our own mentality, but we can transcend the mentality of others; and of course they may transcend ours too. Therefore perhaps together we can transcend ourselves.

We can use truth to explain the success of actions without fear of circularity. The obvious problem we are then left with is to determine precisely what we mean by 'success'. And the solution to this, as we see in the remaining chapters, depends on Darwin.