

Chapter Nine

THE LIMITS OF THINKING WITHOUT WORDS

In the previous chapter we saw that a certain type of thinking is unavailable at the non-linguistic level. A forms of thinking which involve thinking about thought (and hence taking particular thoughts as the objects of thought) are only available to creatures participating in a public language. Thoughts can only be the objects of further thoughts if they have suitable vehicles and the only possible vehicles are public language sentences. Intentional ascent requires semantic ascent. In this chapter I will consider the practical implications which this has for the scope and limits of non-linguistic thought. We will be concerned in particular with the following two questions. What types of thinking are in principle unavailable to non-linguistic creatures? What sort of primitive precursors might there be at the non-linguistic level for types of thinking which involve intentional ascent in their full-fledged form?

It will turn out that intentional ascent is a rather broader category than immediately appears. There are two types of intentional ascent, which might be termed explicit and implicit intentional ascent respectively. The first type includes all those forms of thought involving what is frequently called *metarepresentation*. Metarepresentation involves the ability to have thoughts that take further thoughts as their direct objects. We will consider examples of metarepresentational thinking in the first three sections of the chapter. In 9.1 I explain the distinctive type of belief revision available only to language-using creatures, contrasting it with the more primitive ways in which non-linguistic creatures can modify their beliefs and behavior in the light of changing evidence, whether supporting or countervailing. Section 9.2 discusses how the possibility of a certain type of higher-order desire involves intentional ascent (the

taking of a cognitive attitude towards a particular desire) and hence depends upon language. In section 9. I discuss the relation between language possession and what is often called theory of mind. I argue that the attribution of beliefs, certain types of desires (what I earlier called situation-desires) and propositional attitudes in general is only available to linguistic creatures – again because it involves thinking about thinking. A primitive type of psychological explanation, involving the attribution of goal-desires and perceptual states, is however possible at the non-linguistic level.

Metarepresentational thinking, however, is not the only form of thinking that involves intentional ascent. Intentional ascent can also occur implicitly. Implicit intentional ascent does not involve taking thoughts as the objects of further thoughts. Rather, it involves embedding thoughts to create more complex thoughts whose contents are states of the world, rather than my own mental states, as when one engages in second-order cognitive dynamics, or someone else's mental states, as when one is attributing beliefs and desires to others. In section 9.4 I will explain in the abstract how and why embedding requires intentional ascent. The remaining sections of the chapter will explore different types of embedding and the restrictions thereby imposed upon non-linguistic thinking. Section 9.5 shows how a certain type of narrative thinking is dependent upon language. I suggest that tense operators are best viewed as analogous to adverbial modifiers, qualifying complete thoughts, with the result that tense-based thinking involves a form of intentional ascent. Sections 9.6 and 9.7 explore the distinctive contribution to thought of the relative pronoun, already discussed in Chapter Four as the linguistic analogue of the bound variable of quantification. In 9.6 I discuss the role which the acquisition of the relative pronoun might play in making possible the move from domain-specific cognition to domain-general cognition, while section 9.7 explores the different ways of thinking about generality available at the linguistic and non-linguistic levels.

9.1 *Belief revision and second-order cognitive dynamics*

The distinctiveness of linguistic cognition was introduced in the previous section through Clark's notion of second-order cognitive dynamics. Here again is how Clark characterizes this notion:

By second-order cognitive dynamics I mean a cluster of powerful capacities involving self-evaluation, self-criticism and finely honed remedial responses. Examples would include: recognizing a flaw in our own plan or argument, and dedicating further cognitive efforts to fixing it; reflecting on the unreliability of our own initial judgements in certain types of situation and proceeding with special caution as a result; coming to see why we reached a particular conclusion by appreciating the logical transitions in our own thought; thinking about the conditions under which we think best and trying to bring them about.

In this section I will delineate more precisely the scope of the distinctive type of reflection made available by the presence of language. It will turn out that this distinctive type of reflection has two broad components – the availability of a distinctively second-order variety of belief revision and the availability of second-order desires.

In order to appreciate what is really distinctive about linguistic cognition we will need to make some basic distinctions. The first concerns the relation between behavior and its effects. It is clear that many types of behavior (indeed, plausibly, all types of behavior that are not tropistic, driven by innate releasing mechanisms or classically conditioned) are driven by what might broadly be termed outcome-sensitivity. Animals and infants will modify all but the simplest behaviors according to their outcomes. This is, of course, the essence of trial-and-error learning, as well as being an indispensable adaptive trait. But it is important to distinguish two different types of outcome-sensitivity. We can term them practical outcome-sensitivity and doxastic outcome-sensitivity respectively. A creature is practically outcome sensitive whe

it modifies its behavior directly as a function of exposure to evidence of that behavior's success or failure. A creature is doxastically outcome-sensitive, on the other hand, when it modifies its behavior as a function of modifications in its beliefs, having modified its beliefs in response to evidence of the truth or falsity of those beliefs.

Practical outcome-sensitivity can be illustrated through the distinction between classical conditioning and instrumental conditioning. In classical conditioning a neutral stimulus (e.g. the sound of a bell) is followed by a reinforcing stimulus (e.g. the presentation of food) that elicits a reaction (e.g. salivation). The outcome of classical conditioning is that the conditioned response (the salivation) comes to be given to the conditioned stimulus in the absence of the unconditioned stimulus. In instrumental or operant conditioning, the presentation of the reinforcing stimulus is contingent upon the animal making a particular behavioral response (such as pecking a lever). If the behavioral response does not occur the reinforcing stimulus is withheld. Classically conditioned behavior is not outcome-sensitive in any interesting sense, since it is not the behavior that is reinforced. The conditioning is a function of the strength of the association between the conditioned response (e.g. the sound of the bell) and the unconditioned stimulus (e.g. the presentation of food). The process of conditioning is not determined by anything the animal actually does.

Instrumental conditioning, on the other hand, is clearly outcome-sensitive. It is the behavior itself that is reinforced. Nonetheless, most varieties of instrumental conditioning are not doxastically outcome sensitive. The category of instrumentally conditioned behavior is not fully homogenous and, although a case can be made for regarding certain types of instrumental conditioning as resting upon instrumental beliefs (see, e.g., Heyes and Dickinson 1993, Dickinson and Balleine 1993 and the experiments by Rescorla and Skucy discussed below), the vast majority of instrumentally conditioned behaviors are not best interpreted in terms of beliefs about the contingency between behavior and outcome.

Clear examples of doxastic outcome-sensitivity come with behaviors exemplifying what in Chapter Five I termed level-2 rationality. It will be recalled that it is constitutive of the type of consequentialist thinking characteristic of level-2 rationality that a creature be sensitive to whether the instrumental contingency it is tracking holds or not. Level-2 rationality involves instrumental beliefs about the contingencies between courses of action and outcomes and no creature can properly be credited with such instrumental beliefs unless those beliefs are evidentially sensitive to the contingency in question. We considered a clear example of this sensitivity in Rescorla and Skucy's lever-pressing experiments. Rats trained to press a lever to obtain food will cease lever-pressing when the schedule is changed so that the food appears irrespective of whether the rat presses the lever or not. They are sensitive to the evidential connection between the action of lever pressing and the appearance of the food and modify their behavior accordingly. Creatures (like Hershberg's chicks who persist in moving towards a food-source that retreats from them at twice they rate at which they approach it) clearly lack this sensitivity. Similarly, a direct sensitivity to the outcomes of action is an integral part of tool construction and use. One cannot devise tools without being sensitive to their aptness or not for the job for which they are being devised. Supporting and/or countervailing evidence needs to be taken into account in modifying the tools.

As the example of level-2 rationality shows, doxastic outcome-sensitivity is not simply a matter of behavior modification resulting from changes in beliefs. If, for example, a creature stops stalking one kind of prey because it has noticed an easier quarry within closer reach then its behavior has been modified as a consequence of changes in its (perceptual) beliefs. Nonetheless, it has not displayed doxastic outcome-sensitivity because the beliefs in question are not beliefs about the outcome of the behavior and the means by which that outcome is to be achieved. Nor are the relevant changes in belief the result of changes in the balance of evidence for the original belief (as they would be if, for example, the creature noticed that what it had taken for one kind of prey was really another). Bearing this in mind, many forms of behavior exemplifying level-1 rationality are going to come out as practically but not doxastically outcome-

sensitive. We saw in Chapter Six that a wedge needs to be driven between the normative and descriptive dimensions of our theories of practical decision-making as applied to behaviors falling within the general category of level-1 rationality. We as observers can model level-1 choice behavior in terms of some normative theory such as expected utility theory, which has a manifestly instrumental dimension (given that the calculation of expected utility rests upon the relevant utilities and probabilities of the different possible outcomes consequent upon any given behavioral option). But this does not mean that the practical decision-making resulting in the behavior itself involved any instrumental reasoning – nor, *a fortiori*, any degree of doxastic outcome-sensitivity. A clear example of this emerged in the red deer roaring contexts. The louder a red deer stag roars the stronger it is – since the same thoracic muscles are employed in roaring and fighting. To the extent that a deer acts because it recognizes that it is roaring louder than its antagonist, then, it will be acting in a way that can be analyzed instrumentally as maximizing its expected utility. But it is not acting in the way it does because it has engaged in any consequentialist reasoning. Therefore, even if during the roaring contest the stag discovers that its antagonist can in fact roar louder than initially appeared and consequently withdraws from the roaring contest this would still not count as doxastic outcome-sensitivity.

It is clear, therefore, that doxastic outcome-sensitivity, although relatively circumscribed within the non-linguistic realm, is not the preserve of language-using creatures. To see what is epistemically distinctive about the realm of the linguistic we need to make a further distinction between two different types of doxastic outcome-sensitivity. Doxastic outcome-sensitivity involves, as we have seen, behavior modification as a consequence of doxastic modification – and not simply doxastic modification *per se* but rather the modification of beliefs about the instrumental connection between a particular course of action and an intended outcome. There are two different ways in which the relevant doxastic modification can take place. The examples we have been considering up to now are of what might be termed *direct doxastic modification*. That is to say, beliefs are modified in direct response to changes in the structure

and nature of available evidence. When it becomes apparent to the rat, for example, that the food reward will appear irrespective of whether or not it presses the lever, this new piece of evidence has a direct effect on the rat's belief about the dependence of food upon lever-pressing and that original belief is accordingly modified. A similar process might occur during the process of tool construction – when the stone fractures one way rather than another the experienced stone-knapper might revise their instrumental beliefs about the appropriate way to detach flakes to construct a hand-axe. What is characteristic of both these examples is that perceptually registered countervailing evidence impacts immediately upon the belief in question. The instrumental belief controlling the behavior is instantly modified in response to the perceived discrepancy between actual outcome and envisaged outcome.

In direct doxastic outcome-sensitivity, therefore, there is no need for the rat or the tool-maker to reflect on the relations of evidential support that hold between what they perceive and what they believe. Such reflective changes of belief are, however, characteristic of the second type of doxastic modification (that we can term *reflective doxastic modification*). The distinction can be put in terms of the different norms governing each type of belief revision. Direct doxastic modification can be seen as a process of maintaining and/or restoring coherence within the belief system (and, of course, between the belief system and the deliverances of perception). What matters is that the belief system should be coherent and integrated enough for the creature in question to present a unified epistemic perspective on the world and to be able to act accordingly. There should be no tension (more realistically: as little tension as possible) between the evidence that is available to a creature and the beliefs that that evidence supports. When there is tension between evidence and beliefs modifications are made to restore coherence – and, as we have seen, these modifications are made directly.

Reflective doxastic modification, on the other hand, is governed by the norm of truth rather than the norm of coherence. Coherence is, of course, a desideratum but it is not the sole desideratum. What

matters above all is that the evidence should be such as make the beliefs that it supports likely to be true. Reflective doxastic modification concerns itself explicitly with the logical and probabilistic relations between evidence and beliefs (as well, of course, as between the individual beliefs within a belief system). At the reflective level a belief might be rejected or modified in the absence of countervailing evidence or tension with existing beliefs – it might, for example, be rejected simply because the believer recognizes that it is not warranted by the evidence. Although the notions of warrant and justification can be applied in an attenuated sense at the level of direct belief revision (most prominently in a broadly externalist sense according to which, roughly speaking, a set of beliefs is warranted to the extent that it is produced by reliable mechanisms and modified according to principles that tend to preserve truth and eliminate error), there is no sense in which the extent to which their beliefs are warranted or justified can be an issue for creatures operating solely at that level. That is to say, internalist notions of warrant and justification can get a grip only at the reflective level, because these notions can be applied only to thinkers capable of explicit reflection on the relations between thoughts and perceptions and between thoughts and other thoughts.¹

The link between reflective doxastic modification and language should be clear, given the discussion in the previous section. It is only possible to evaluate and reflect upon the extent to which one belief implies another (more precisely: the extent to which acceptance of one belief commits one to acceptance

¹ The distinction between direct and reflective doxastic modification is linked to the distinction drawn by Gilbert Harman between foundationalist and coherentist modes of belief revision (see Harman 1986). Harman's book contains much useful discussion of the different principles operating in the two different modes of belief revision. He does not take a view, however, on the relation between the coherentist/foundationalist distinction and the divide between the linguistic and the non-linguistic.

of another), or the extent to which a belief is supported by a particular type of evidence if one is able explicitly to hold those beliefs in mind. And, as we saw at some length earlier, thoughts can only be the objects of further thoughts in the way that this requires if they have linguistic vehicles. It is clear, moreover, that reflective belief revision will involve explicit consideration of the formal logical relations between thoughts – and we have no understanding of the logical relations between thoughts except when those thoughts have linguistic vehicles.

The discussion so far has concentrated on processes of belief revision. But once the mechanisms of reflective belief revision are in place broader possibilities emerge for reflective self-monitoring and epistemic self-criticism. Reflective subjects can evaluate beliefs without changing them. They can distinguish between those beliefs that are candidates for revision and those that are not candidates for revision. They can plot out the evidential connections between the different components of their belief systems, working out what rests on what and which beliefs are bearing the greatest weight. Reflective subjects can also identify and learn from their past mistakes, settling on epistemic policies that are designed to avoid those mistakes. The notion of second-order cognitive dynamics that Clark originally introduced includes all these high-level forms of epistemic self-monitoring. They are all rooted, however in the distinctive type of belief revision that I have termed reflective doxastic modification – itself a particular type of outcome-sensitivity. The process of emancipation from the doxastic here-and-now is a gradual one.

9.2 *Higher-order desires*

Non-linguistic creatures are incapable of intentional ascent. We have seen the role that intentional ascent plays in certain types of belief revision and, more broadly, in epistemic self-monitoring. Are there

comparable phenomena in the realm of desire? It is useful at this point to advert to a distinction again made some years ago by Harry Frankfurt in an influential discussion of freedom of the will and the nature of personhood (Frankfurt 1971). Frankfurt distinguishes first-order from second-order desires. In the terminology we have been adopting first-order desires can be either goal-desires or situation-desires. They are first-order because they are directed at objects and/or properties (in the case of goal-desires) or states of affairs (in the case of situation-desires). Second-order desires, in contrast, are situation-desires that take first-order desires as their objects. A second-order desire is a desire to the effect that one should have a particular first-order desire.²

When we consider the structure of second-order desires it becomes clear why an important class of second-order desires should be language-dependent. The content (that is to say, the satisfaction condition) of a second-order desire is that one should have a particular first-order desire. The content of the first-order desire is embedded within the content of the second-order desire. It is, so to speak, the target of the second-order desire. But then it looks very much as if second-order desires involve intentional ascent and therefore, by the argument of section 7.2, the capacity for semantic ascent. Still, some qualification is required. In particular we need to distinguish cases in which the target first-order desire is a goal-desire from cases in which it is a situation-desire. The satisfaction condition for a second-order desire embedding a goal-desire is the state of affairs in which the desirer stands in an extensionally specifiable desiring relation to a particular object or property. This does not seem to involve intentional ascent in the strict sense of the word, since the content of the target desire is specified only at the level of reference.

² Frankfurt makes a further distinction between second-order desires and second-order volitions, where a second-order volition is a desire not simply to have a particular first-order desire but also for that first-order desire to be one's will. I will use the term "second-order desire" to include second-order volitions.

There is no embedded thought, which there clearly is when the target desire is a situation-desire. The satisfaction condition of a situation-desire is the state of affairs in which the desirer stands in the desiring relation to the state of affairs in which he has a desire with a determinate structured content at the level of sense.³

Of course, the capacity to form second-order desires is not an isolated capacity. Just as reflective belief revision comes as part of a package of strategies and abilities for epistemic self-assessment and self-monitoring, so too should the capacity for second-order desires be viewed as part of a complex set of strategies and abilities for determining the sort of person one is going to be. The second-order desirer is capable not simply of assessing his first-order desires, but also of evaluating his first-order desires one by one and indeed of comparing and weighting them. New possibilities of decision-making open up. The second-order desirer can reflect on desires and preferences in a way that will allow him to adopt something like the primitive version of decision theory that we considered at the beginning of Chapter Six. One would expect, moreover, that this type of reflection would be linked with a particular type of self-awareness – the awareness of oneself as an agent with distinctive projects, goals and aims. The agent who can reflect on his own desires will also be able to see how they fit together, which ones are more fundamental and which are more basic to his character. The analogy with the epistemic case is very close. In each category the basic capacity for intentional ascent is linked with reflective abilities that have no analogue at the non-linguistic level.

³ As was noted in Chapter Three, this type of second-order desire can be accommodated within the general paradigm of success semantics. I suggested there that those second-order desires which are targeted on specific first-order desires can be understood as states which cease when the relevant first-order desires come into being.

9.3 *Intentional ascent and understanding other minds*

It is natural to think that the conclusions of the previous two sections can be extended from the intrasubjective case to the intersubjective case. We have seen that intentional ascent, and therefore language, is required for reflection on one's own beliefs and desires. Might a parallel argument hold for reflection on the beliefs and desires of other subjects and agents? The important question, though, is how far the parallel argument can be applied. Is it possible for non-linguistic creatures to participate in practices of attributing psychological states to their conspecifics or indeed to any other creatures? We have seen that non-linguistic creatures can have a surprisingly rich psychological life. But to what extent can they understand the psychological lives of others? Unsurprisingly it turns out that everything depends on which psychological states are in question. In this section I shall argue that only a very limited range of psychological attributions is possible at the non-linguistic level. In essence, all that is available at the non-linguistic level is an understanding of the perceptual states of other creatures, an understanding of their goal-desires and, relatedly, an understanding of the goal-directed nature of intentional action.

In the light of the preceding discussion it is not hard to see why a very broad class of psychological attributions should be unavailable to non-linguistic creatures. To attribute, for example, a belief or a situation-desire to another creature is essentially to view that creature as standing in a particular relation to a thought – the relation of believing the thought to be true or the relation of desiring that the state of affairs characterized in the thought come to pass. Clearly, therefore, the attribution of a belief requires thinking about a thought. It is a canonical form of intentional ascent that requires being able to “hold a thought in mind”. As such it is immediately susceptible to the argument sketched out in section 7.2 to the effect that intentional ascent requires the possibility of semantic ascent and hence is only available to

language-using creatures. A belief can only be attributed if the thought that is the content of the belief can be represented by the attributer, which requires that it be vehicle – and, as we have seen, the only candidate vehicles are linguistic.

The argument applies to the attribution of all mental states that are traditionally known as propositional attitudes – that is, to the attribution of all those mental states that involve the subject's taking a particular attitude (the attitude of hoping, say, or fearing) to a thought. This is closely connected of course, with two familiar facts about ascribing propositional attitudes. The first is that propositional attitude ascriptions create opaque contents – they are intensional. A creature may believe a particular state of affairs to be the case under one “mode of presentation” (under one way of characterizing it) but not under another. Opacity is one fundamental reason why we need to take the objects of propositional attitudes to be thoughts, rather than the states of affairs or situations that are the truth-conditions of those thoughts. The second familiar fact (and the second fundamental reason for the opacity of propositional attitude ascriptions) is the possibility of error. One can believe things that are not true and one can desire things that do not exist (that, for example, will not come into existence unless the desire is satisfied). How can the object of belief be a situation or state of affairs if it is possible to believe things that are not in fact the case? What could the object of belief be other than a thought? For these two reasons, then, there can be no attribution of propositional attitudes without intentional ascent.

So, any mental state that involves taking an attitude to a thought can only be attributed by language-using creatures (although of course it can perfectly well be *had* by non-linguistic creatures). If there are any mental states that can be attributed by non-linguistic creatures, therefore, they can only be mental states that do not involve taking an attitude to a thought – and that, in turn, means that they can neither be opaque nor admit the same type of possible error that we find in the canonical propositional attitudes. It is easy to identify one candidate on the basis of the distinction between goal-desires and situation-desires

made in Chapter Two and reiterated in the previous section. The distinction is effectively between desire construed as a propositional attitude (in situation-desires, that are attributed via that-clauses that pick out that thought that is the object of desire) and the more fundamental goal-desires that are directed not at thoughts but rather at objects or features. It will be recalled that goal-desires, unlike situation-desires, are extensional. There is no reason, therefore, why non-linguistic creatures should not be able to attribute goal-desires to other agents. The argument from intentional ascent cannot get a grip, since goal-desires are relations between a subject and an object/feature, rather than between a subject and a proposition.

The ability to attribute goal-desires goes hand in hand with a basic understanding of intentional, that is to say goal-directed, behavior. Although of course there will be many different degrees of complexity in goal-directed behavior, depending on the richness of the desires and beliefs by which it is driven, a creature that is capable of attributing goal-desires will be able to make the basic distinction between purposeful behaviors, on the one hand, and random movements and instinctive reactions on the other. A purposive action is an action for which a motivating goal-desire can be identified.

Goal-desires cannot be the only mental states that can be identified and attributed by non-linguistic creatures. It is hard to see, for example, how a goal-desire can be attributed to a creature without some evidence of the information that the creature possesses about its environment. At the bare minimum this information will be perceptual. To know what goal-desire might be motivating a creature at a given moment we need to know, first, what end it is pursuing and, second, how it might reasonably expect that end to be realized by its current behavior. Both of these require knowing to which features of its environment the creature is perceptually sensitive. If, therefore, a non-linguistic creature is to be able to attribute goal-desires to a fellow creature it must be able to formulate hypotheses about what that creature is perceiving.

This presents us with a puzzle. Reports of perception are usually taken to bear significant structural commonalities with propositional attitude reports. Except for those philosophers adopting some version of the epistemic theory of perception (according to which perception can be reductively analyzed in terms of the acquisition of dispositions to belief), there is a consensus that the content of perception cannot straightforwardly be assimilated to the content of belief (Dretske 1981, Peacocke 1983, Peacocke 1992, Crane 1992). Nonetheless, the acknowledged differences between the content of belief and the content of perception are not usually taken to stand in the way of reporting the content of perception with “that-“ clauses specifying the perceiver’s relation to a thought – the thought in question being the propositional content of the belief to which the perception might be expected to give rise.⁴ But this, by the familiar argument from intentional ascent, seems to place the understanding of perception beyond the reach of non-linguistic creatures – and therefore (by the argument of the previous paragraph) bars them from identifying and attributing goal-desires.⁵

⁴ This is perfectly compatible with rejection of the epistemic theory of perception. The proposal is not that the propositional report says all that there is to say about the content of perception. One might think that perceptual states carry information in an analogue rather than digital manner, for example, while nonetheless thinking that the information carried can be propositionally reported. Nor, of course, does the suggestion that the propositional content of perception corresponds to the content of the perceptual belief to which it might be expected to give rise imply that perceptions are nothing more than dispositions to acquire belief. The canonical statement of the epistemic theory is Armstrong 1968.

⁵ It is worth noting that, strictly speaking, perceptual reports only can only have one of the two characteristic features of propositional attitude reports. Clearly, the “that-“ clauses reporting the contents of perception are

The resolution of the puzzle comes with a distinction between two different types of perception and, correspondingly, between two different types of perceptual report. I will call the two types of report SS-reports and ES-reports respectively.⁶ SS-reports have the following three characteristics. They

- (i) Take non-propositional complements
- (ii) Create transparent contexts
- (iii) Are made true by direct perceptual relations between perceivers and particulars

ES-reports differ from SS-reports in all three dimensions. ES-reports

- (i) Take propositional complements

intensional, but since perception is factive (that is, one can only perceive what is in fact the case), perceptual reports cannot be false. All this means, of course, is that we sometimes have to refer to “seemings-to-perceive” rather than to perceivings tout court.

⁶ The basic distinction here is made in various places, including Dretske 1969 and the final chapter of Jackson 1975. In the characterisation of SS-reports and ES-reports I am drawing on a more recent paper by Kevin Mulligan (Mulligan 1998).

(ii) Create opaque contexts

(iii) Are made true by perceptually-based epistemic relations between perceivers and propositions

Many philosophers endorse the idea that SS-reports and ES-reports characterize different modes of perception. In Dretske's terms, SS-reports describe what can be termed *simple seeing*, while ES-reports characterize *epistemic seeing*.⁷ The basic idea is that simple seeing and epistemic seeing are dissociable components in normal visual perception. Dretske, in the book that put this sort of distinction on the map, explicitly suggests that what we see in simple seeing (or what he calls non-epistemic seeing) "is a function solely of what there is to see and what, given our visual apparatus and the conditions in which we employ it, we are capable of visually differentiating" (1969, p.76). That is to say, the distinction between simple seeing and epistemic seeing is phenomenologically salient. Dretske explicitly draws a parallel with the conception of the sensory given that was at the core of the sense-datum theory of perception. Whereas the sense-datum theorist believed, or is alleged to have believed, that we can strip away from the rich manifold of perceptual experience to reach a level at which what is really seen are colored expanses, the simple seeing theorist holds that what we really reach after such a process is correctly characterized by an SS-report.

For present purposes we can remain neutral on the question of whether this dissociability thesis is correct at the level of perception itself. It may or may not be the case that we can strip away the epistemic dimension from the perceptions of creatures capable of epistemic seeing in a way that will allow us to characterize these perceptions entirely in non-epistemic terms. It is certainly true, however, that the

⁷ As is frequently the case in the philosophy of perception the modality of vision is really what's under discussion.

dissociability thesis is perfectly correct at the level of perceptual reports. An episode of epistemic seeing can perfectly well be characterized, albeit incompletely, by an SS-report. The SS-report can be viewed as specifying the objects of perception - I am taking the objects of perception here to be whatever it is that stands at the other end of the causal chain that leads to perceptual states. On this report, an SS-report simply states that the perceiver stands in an appropriate perceptual relation to a certain object – the relation that will allow him or her to discriminate that object from the perceptual background. It has nothing to say about how that object is perceived (about its mode of presentation), which of course is why SS-reports create transparent contexts.

The significance of this should be clear. The argument from intentional ascent shows that non-linguistic creatures are not capable of the sort of understanding of vision that might be expressed through ES-reports, since ES-reports require specifying a perceiver's relation to a thought. But this is perfectly compatible with non-linguistic creatures being capable of SS-reports and, correlatively, of thinking about the direct perceptual relations in which other creatures stand to objects. And nor are SS-reports, and the understanding that goes with them, quite as restricted as initially appear. Although many of those who have made the distinction between simple seeing and epistemic seeing have specified the former in terms of relations purely between a perceiver and an object, there is no reason why this has to be the case. The *relatum* of simple seeing can be a state of affairs, where this might minimally be construed as an object exemplifying a property or two or more objects exemplifying a particular relation.⁸ It seems natural to suggest that a non-linguistic creature can include in the complement of an SS-report pretty much anything

⁸ Mulligan suggests that the characterisation of simple seeing should employ the machinery of tropes rather than properties (Mulligan 1998).

it can think about directly – that is, anything it can think about directly it can think about as the object of another’s perceptual state.

The availability of SS-reports to non-linguistic creatures allows them to engage in a primitive form of psychological explanation. Psychological explanation, as traditionally construed, is not available to non-linguistic creatures since it depends upon the attribution of beliefs and situation-desires. Nonetheless, a combination of SS-reports and the attribution of goal-desires will permit non-linguistic creatures to make some progress on explaining and predicting the behavior of their conspecifics and potential predators. If one knows what a creature desires and has some sense of its perceptual sensitivity to the environmental layout (as well of course as an understanding of its motor capabilities) one can expect to be able to predict its behavior with some success.

To take stock, then, it appears that propositional attitude ascriptions are beyond the capacity of non-linguistic creatures. The argument from intentional ascent shows us that language is required for all psychological attributions that specify a thinker’s relation to a thought. The only psychological attributions that escape the argument from intentional ascent are attributions of goal-desires and the SS-reports that characterize simple seeings (and the equivalent, of course, in other sensory modalities). Non-linguistic creatures, therefore, are capable only of a highly circumscribed understanding of other minds that falls far short of what is often described as possession of a theory of mind.

It should be noted that this restrictive interpretation of the “mind-reading” abilities of non-linguistic creatures is compatible with much recent research into the extent to which non-human primates can properly be described as possessing a “theory of mind”. There are well-documented examples of primate behavior that some prominent students of animal behavior have thought can only be interpreted as examples of interpersonal deception (Premack and Woodruff 1978, de Waal 1982 and some of the essays

in Byrne and Whiten 1988). But the consensus opinion among primatologists is that a more parsimonious interpretation of these behaviors is to be preferred (see, e.g. Gómez 1996, Povinelli 1996 and Hauser 2000).⁹ Many of the examples of what has come to be termed *tactical deception* (Byrne 1995) can be understood as the manipulation, not of another's propositional attitudes, but simply of their visual perspective. Here is an example of a tactical deception in a troupe of baboons in Ethiopia that lends itself to such an interpretation:

An adult female spent 20 min in gradually shifting in a seated position over a distance of about 2m to a place behind a rock about 50 cm high where she began to groom the subadult male follower of the group – an interaction not tolerated by the adult male. As I was observing from a cliff slightly above [the animals] I could judge that the adult male leader could, from his resting position, see the tail, back and crown of the female's head, but not her front, arms and face: the subadult male sat in a bent position while being groomed, and was also invisible to the leader. The leader could thus see that she was present, but probably not that she groomed. (Report by Hans Kummer quoted in Byrne 1995 p. 106)

The behavior of the female baboon, assuming that it is indeed to count as an instance of tactical deception, does not seem to require assuming an intention to manipulate the beliefs of the alpha male (e.g. an intention to bring it about that he believe that she is not grooming the subadult male). What she is

⁹ Michael Tomasello has explicitly argued that (to put it in my terms) the capacity for intentional ascent is unique to humans and hence that the mind-reading interpretation of non-human apes cannot be correct (Tomasello 1998). He does not, however, offer an argument for why intentional ascent should not be available non-humans. The argument from intentional ascent to semantic ascent would no doubt be congenial to him.

doing is profiting from an understanding of the alpha male's visual "take" on the situation to escape detection. The female baboon needs only to appreciate the alpha male's line of sight and the fact that he would be prevented from seeing the subadult male by the intervening rock. This seems firmly at the level of simple seeing rather than epistemic seeing.

It is interesting to note that a recent survey article draws a basic distinction between two different levels of understanding vision that can be mapped without too much effort onto the distinction between simple seeing and epistemic seeing and that can be put to work to distinguish between different ways of understanding the experimental and ethological evidence for primate "mind-reading". Daniel Povinelli has distinguished three different types of knowledge about visual perception and its mechanisms (Povinelli 1996). The first level is a simple sensitivity to the presence of eyes and eye-like stimuli. This sensitivity is exploited, for example, by the many species of moth and butterfly that are patterned to resemble eyes. Such sensitivity is clearly widespread throughout the animal kingdom, and is too primitive to concern us. The second level seems closely correlated with an appreciation of non-epistemic seeing. Povinelli describes it as an understanding of perception as a cognitive connection between organisms and world. The third level seems much closer to an understanding of epistemic seeing. This is "the understanding that in addition to linking an individual's mental state of attention to the external world, visual perception also alters one's internal experiences, states of knowledge and belief" (Povinelli 1996, p.313). Within this classification the standard understanding of primate social deception would place it squarely at the third level, while the proposal being developed in this section is that it can be understood at the second level.

Considerable laboratory research has recently been carried out on primate understanding of visual perspective and the direction of gaze. Some of this research has been negative. There is solid experimental evidence that chimpanzees are not very good at understanding the visual perspective of

humans (Povinelli and Eddy 1996a, 1996b). In particular, when given the choice of begging for food from an experimenter who can see them from one who cannot, they are only capable of making very crude discriminations. They successfully distinguish between, for example, an experimenter facing them and one whose back is turned towards them, but they do not seem to understand the difference between an experimenter who is directly looking at them and one who cannot see them because he has a bucket over their head. Some experimenters have taken this to cast doubt upon chimpanzee understanding of vision and visual perspective. Others have pointed out, however, that the paradigm is a highly unnatural one for primates, not simply because it involves humans rather than conspecifics but because the natural situation for chimpanzees is one of competition for scarce resources, rather than cooperation or collaboration (Hauser et al. 1993). Recent investigation of chimpanzees' understanding of the visual perspective of conspecifics in competitive situations shows that chimpanzees can monitor the visual perspective of a competitor and information to guide their own actions (Hare et al. 2000, Hare, Call and Tomasello 2001). One group of experiments involved pairs of chimpanzees, one subordinate and one dominant, placed in separate rooms and separated by occluders (Hare, Call and Tomasello 2001). Experiments hid food in full view of the subordinate chimpanzee that could expect that the dominant chimpanzee would be successful in competing for any food which both had seen hidden. The various conditions manipulated the dominant chimpanzee's visual access to the location of the food – in some he saw where the food was hidden, in others not, while in a third condition he was misinformed about the location of the food because it was moved after he had initially seen it being placed. In all conditions the subordinate was able to monitor the dominant's visual access to the location of the food. The subordinate chimpanzee was consistently able to make use of this information to its own advantage, refraining from directly competing with the dominant chimpanzee for food about whose location the dominant had accurate information and preferentially retrieving food items about which the dominant was either uninformed or misinformed.

A natural way of interpreting these experimental data, in conjunction with the ethological data which originally gave rise to them, is as manifesting a form of social understanding which relies upon the use of information about vision and goal-directed action, rather than upon the attribution of beliefs, desires and other propositional attitudes. If the argument of this chapter is correct then this form of social understanding is all that is available at the non-linguistic level and is sufficient to explain the various forms of social deception, social interaction and social cognition which have been identified in non-linguistic creatures. This proposal is in line with that offered by primatologists such as Tomasello and Call (1997) and Whiten (1996) who complain that the existing discussions of primate social cognition are unnecessarily polarized between interpretations of primate social behavior in terms of full-fledged belief-desire psychology, on the one hand, and unthinking reinforcement on the other.

9.4 *Embedding and intentional ascent*

The role played by intentional ascent is very clear in the three types of thinking we have considered. Thoughts are part of what is being thought about in second-order cognitive dynamics, in second-order desires and when one speculates about the mental states of others. When, for example, one attributes a belief to another creature one is effectively attributing to that person a relation to a thought. Similarly, in evaluating a piece of reasoning in the manner characteristic of second-order cognitive dynamics one is considering the logical or probabilistic relations holding between the thoughts of which it is made up. It does not make sense to suppose that either of these types of thinking could be engaged in without the capacity for intentional ascent – and hence, by the argument of the previous chapter, without the possibility of semantic ascent. It would be a mistake, however, to conclude that intentional ascent is involved only when thoughts are explicitly targeted on other thoughts – that is to say, when thoughts are

part of the content of what is being thought. I will bring out in this section how the capacity for intentional ascent is required for all thinking that involves compound thoughts with further thoughts embedded in them, irrespective of whether those embedded thoughts feature in the content of thinking.

We can start with a basic class of compound thoughts – namely, those involving the basic logical connectives, such as disjunction, conjunction and the material conditional. Consider a disjunctive thought of the sort that might be expressed in the sentence ‘A or B’. What is it to be capable of entertaining such thought? It is to be capable of understanding that a certain relation holds between two thoughts – the relation of their not both being false (and, on some understandings of ‘or’, their not both being true). Of course, the disjunctive thought itself is not a thought about the two thoughts expressed by ‘A’ and ‘B’. It is a thought about the states of affairs that serve as the truth conditions for the thoughts. Nonetheless, the disjunctive thought is not available to be thought by any creature that is not capable of thinking about how the truth-value of one thought might be related to the truth-value of another thought.¹⁰ The same holds of the other truth-functional propositional operators. Understanding a truth-functional operator is understanding how it serves to form compound thoughts whose truth-value is a function of the truth-values of their parts. No creature that was not capable of thinking about thoughts could have any understanding of truth-functional compound thoughts.¹¹

¹⁰ As discussed in Chapter Seven, analogues of disjunction are available at the non-linguistic level – but these analogues do not employ anything like operators on propositions.

¹¹ Of course, there are philosophers who think that the primary bearers of truth and falsity are sentences rather than thoughts. This would provide an even quicker way of arguing from truth-functional connectives to the need for semantic ascent.

The example of the truth-functional propositional operators shows that a particular type of thinking can presuppose the capacity for intentional ascent, even though the thoughts instantiating that type of thinking are not themselves thoughts about thoughts. Compound thoughts formed through the truth-functional propositional operators involve intentional ascent *implicitly* rather than explicitly. Are there other types of complex thoughts of which the same can be said? Extrapolating from the truth-functional operators suggests that a similar argument will apply to complex thoughts formed by operators that can only be understood via the notion of truth – and this, of course, can hold for operators that are not truth-functional. Let me offer a conjecture with respect to modal operators (in the sense of ‘modal’ on which modal operators are those attributing possibility and necessity). It seems very plausible that our understanding of necessity cannot be divorced from the understanding of truth. Our primary understanding of the notion of necessity may well be in terms of a propositional (or, for that matter, a sentential) operator, and a proposition is necessary just if it is true in every possible situation (in every possible world). If this is the case then modal thoughts would presuppose intentional ascent, and with it semantic ascent, no less than compound thoughts formed by means of the truth-functional propositional connectives.

Modal thinking is in many respects similar to tensed thinking. The logical systems that have been developed to display the logical structure of tensed thought are analogous to those that have been developed to elucidate the logical structure of modal thought. They are all what might be termed adverbial extensions of non-modal and non-tensed logics.¹² That is to say, they can be derived from non-

¹² I am taking ‘extension’ in a non-technical sense here. But of course both modal and tense logics are extensions of classical logic in the technical sense that they incorporate the vocabulary of classical logic and have the same theorems and inferences involving *only* that vocabulary.

modal and non-tensed logics by the application of operators that function as adverbs.¹³ The following passage is a classic statement of this approach as applied to tense logic.

If an expression constructs a sentence out of one other sentence it is an adverb or an adverbial phrase, like 'not' or 'It is not the case that' or 'allegedly' or 'It is alleged that' or 'possibly' or 'It is possible that'. . . I want to suggest that putting a verb into the past or future tense is exactly the same sort of thing as adding an adverb to the sentence. 'I *was* having my breakfast' is related to 'I am having my breakfast' in exactly the same way as 'I am *allegedly* having my breakfast' is related to it, and it is only a historical accident that we generally form the past tense by modifying the present tense, e.g. by changing 'am' to 'was', rather than by tacking on an adverb. In a rationalized language with uniform constructions for similar functions we would form the past tense by prefixing to a given sentence the phrase 'It was the case that' or 'It has been the case that' and the future tense by prefixing 'It will be the case that'. (Prior 1968, p.7 quoted in Evans 1985)

The central implication of the adverbial nature of the tense and modal operators is that, where 'T' is an operator of the relevant class and 'p' the name of a sentence, the meaning of 'Tp' will be a function of the meaning of 'p' and the meaning of the relevant operator.

Perhaps it is equally the case that our understanding of temporal order is inextricably linked with our understanding of truth, so that understanding tense is a matter of understanding that propositions can have different truth-values at different times. Our basic understanding of temporal relations is not an

¹³ The general approach has been challenged. See, for example, Evans 1985. But these challenges have not been accompanied by any alternative construals of tense logics.

understanding of temporal relations holding between events or states of affairs. It is, rather, an understanding of the temporal relation between the truth of different propositions. We understand the idea that event A took place earlier than event B by understanding that it was true that event A took place before it was true that event B took place. This is because thinking about the temporal relations between events requires thinking about events holding at particular times, and the understanding of an event holding at a time other than the present is parasitic upon the understanding of a particular proposition being true at a particular time. Thinking about temporal relations (or so I am suggesting) requires being able to think about the possibility of propositions being true or false at different times. It may well be that the capacity to think about possibility, to think about time, and to think about the truth-values of propositions are all interdependent and inter-linked. If this were so, then modal and tensed thinking would be just as unavailable to non-linguistic creatures as thinking involving the truth-functional propositional connectives.

This suggestion about the language-dependence of tensed thinking enables us to make sense of a puzzling transition in the pre-history of the human race. The consensus among archeologists and students of human evolution is almost universal that the crucial stage in human cognitive evolution occurred about 40,000 to 35,000 years ago, with the transition from the Middle Paleolithic to the Upper Paleolithic. This transition seems to have involved a sudden explosion in tool technology and social/cultural organisation, with the emergence for the first time of forms of life that are recognisably congruent with those of modern humans. With the transition to the Upper Paleolithic come the first decorative objects; the first really compelling evidence for totemistic/religious behaviour, as revealed in burial practices and totemic representations; sophisticated hunting strategies that capitalise on seasonal migrations and fluctuations in animal numbers; and far more complex forms of tool production that seem to have drawn upon detailed knowledge of natural history to tailor tools for particular hunting tasks, as opposed to the more general purpose tools of the Lower and Early Paleolithic (for overviews see Donald 1991 Ch. 8, Mellars 1996

and Mitthen 1996a Ch. 9). From this point on the rate of cognitive evolution accelerated exponentially. It is, of course, for this reason that many cognitive archeologists have identified this transition as involving the emergence of a recognizably human language.

One striking feature of the type of cognition emerging in the Upper Paleolithic is the way it exemplifies a novel conception of time. It is a type of thinking reflecting a conception of time that goes beyond a simple ability to discriminate events that are taking place in the present from those that have taken place in the past or that will take place in the future. There are ways of being oriented towards the past or the future that are available to relatively simple organisms – most simply, in the form of expectations about future events and memories of past events. But these need to be sharply distinguished from modes of thought that depend upon thinkers being able to locate themselves within a temporal narrative.¹⁴ What the fossil record tells us about the earlier hominids strongly suggests that they behaved in ways that involved distinguishing the past, the present and the future. But there is nothing to suggest patterns of behaviour available only to creatures capable of narrative thought. In contrast such patterns of behaviour seem extremely widespread in the Upper Paleolithic.

A very clear illustration comes with the hunting strategies that emerged during this period. Archeologists have found a vastly increased sensitivity to seasonal changes in the density and movement of animals. Whereas in the Middle Paleolithic hominids tended to hunt single animals in a very opportunistic way (or indeed perhaps only to scavenge on their remains), the later hominids appear to have shown far more long-range planning in their foraging behavior. They prepared tools long before

¹⁴ My proposal bears certain resemblances to Donald's description of the Upper Paleolithic in terms of mythic culture (Donald 1991). See also the discussion of different ways of understanding time in Campbell 1994.

needing them, designing them for specific prey and specific hunting techniques. Their excavated camp sites seem to have been prepared for long occupations and were systematically moved to intercept prey or harvest food resources (Binford 1989, Mellars 1996, Mitthen 1996a). There is also evidence that, at least after the earlier stages of the Upper Paleolithic, the targets shifted from isolated animals to large herds of animals like reindeer and red deer which were ambushed on their migration paths at particular times of year (Mitthen 1990).

If, as I have suggested, tensed thinking is only available to language-using creatures, then it is easy to see why the evolution of a recognizable human language should have been connected with the emergence of these types of cognition and behavior in the transition from the Middle to the Upper Paleolithic. These forms of hunting behavior presuppose a type of narrative thinking that is quite simply unavailable in the absence of language.

We have so far in this section considered several ways in which intentional ascent might be involved implicitly in types of thinking that are not explicitly targeted on further thoughts. I argued that the formation of compound thoughts by means of the truth-functional propositional connectives presupposes the capacity for intentional ascent and suggested, more tentatively, that similar capacities may be presupposed by the ability to entertain thoughts involving temporal or modal notions. In the remainder of this section I will discuss a further way in which the capacity for intentional ascent can be presupposed by different types of thinking. The types of thinking that we have been considering up to now have involved the formation of complex thoughts by what we can think of as adverbial operators on propositions, such as the truth-functional connectives and the modal and tense operators. Let us turn now to a different way of forming complex thoughts.

Consider the inference form of existential generalization. This is the pattern of inference instantiated by the transition from ‘ Fa ’ to ‘ $\exists xFx$ ’ – that is to say, from an atomic sentence to the effect that a named individual has a given property to the general proposition that at least one individual has that property. The logical operations involved in this transition are clear enough (Dummett 1973). The first is breaking down the atomic sentence into two components, a predicative component and a nominative component (or, in Fregean terms, a function and an argument). Once the internal structure of the atomic sentence is manifest, the next operation is to replace the nominative component by a variable. The final operation is to bind that variable with an existential quantifier. This sequence of logical operations gives us an important clue as to what is involved in a subject’s being able to understand the existential quantifier in a manner that permits existential generalization. In order to understand how a given sentence can imply an existential generalization a thinker needs to be able to view it as being composed in such a way that the nominative component can be replaced by an arbitrary name (and hence by a variable). The sentence needs to be “broken down” in thought before the existential quantifier can be applied. But this breaking down in thought of an atomic sentence presupposes the capacity for intentional ascent. It involves holding the thought in mind and determining its structure in a way that creates a space for the variable that will be bound by the existential quantifier. The same point holds, *mutatis mutandis*, for thinking involving the universal quantifier (as will be discussed further in 9.5 below).

It would seem, then, that there are two principal ways in which the capacity for intentional ascent can be presupposed in the generation of complex thoughts. First, the formation of compound thoughts by means of the truth-functional propositional operators presupposes the capacity for intentional ascent because grasping a complex thought is a matter of grasping the way in which its truth value is determined by the truth-value of its constituent thoughts, and truth is a notion that applies primarily to thoughts. Second, complex thoughts involving quantification presuppose the capacity for intentional ascent because mastery of the mechanisms of quantification requires the ability to determine the structure of a thought. I

suggested also that the capacity to form complex thoughts involving modal and tense operators may well also be dependent upon the capacity for intentional ascent, but this suggestion will play no further role in the argument.

If this line of argument is well-founded then it follows that logic requires language – a conclusion reached on very different grounds by Jonathan Bennett (Bennett 1976). So what are the consequences for reasoning at the non-linguistic level? We have already explored some relevant terrain in Chapter Seven where I showed how non-linguistic creatures could be capable of certain forms of proto-inference that do not involve intentional ascent. Conditional reasoning, it was suggested, can be understood in causal terms while proto-negation can be understood in terms of mastery of contrary concepts (as opposed to the ability to apply a truth-functional operator to propositions). This still leaves open one important question however. If thinking involving quantifiers presupposes the capacity for intentional ascent, and hence is unavailable to non-linguistic creatures, should we conclude that non-linguistic creatures cannot be credited with any form of general thoughts? This question will be addressed in the next section.

The proposed analysis of quantificational thinking raises a further question. I have argued that the key feature in quantificational thinking is the ability to break down a thought in such a way as to open up a space for a variable. Given the highly plausible suggestion, already discussed in Chapter Four, that the variable of quantification is the analogue of the relative pronoun in natural language, it is natural to wonder whether the ability to break down thoughts that comes with the capacity for intentional ascent might not have cognitive implications that we have not yet discussed. In section 9.6 I shall propose that this type of intentional ascent makes possible (and may well be required for) the integration of different types of domain-specific thinking.

9.5 Two types of general thought

In the previous section I argued that quantificational thinking requires language. This conclusion seems to entail that non-linguistic creatures cannot have general thoughts. This section argues that this conclusion only follows on one way of understanding generality.

Many philosophers have distinguished two ways of thinking about generality (Lewis 1969, Bennett 1977). The distinction goes back at least as far as the twelfth century – we find it in Abelard's distinction between generality *in sensu composito* and generality *in sensu diviso*. If I believe, *in sensu composito*, that every F is a G then the content of my belief is a general proposition to the effect that everything that is an F is also a G. There may, of course, be individual things that are F but which I do not believe to be G – since I may be unaware that they are in fact Fs, for example. Nonetheless, as far as I am concerned it is a general truth that nothing can be an F without also being a G. So, for example, I might believe *in sensu composito* that all tigers are dangerous even though, unaware that an albino tiger is still a tiger, I might find myself in front of a tiger and fail to recognize that it is dangerous. If, on the other hand, I believe *in sensu diviso* that every F is a G the content of my belief is not a general proposition at all. Rather, what I have is a certain disposition to form singular beliefs. Whenever I encounter an F, or indeed whenever I think about one, I tend to come to the view that it is a G. It would be evidence against ascribing to me the belief *in sensu diviso* that every tiger is dangerous if I were to fail to identify the albino tiger as dangerous.

The distinction is sometimes termed the distinction between collective generality (*in sensu composito*) and distributive generality (*in sensu diviso*). It can most easily be viewed as a distinction of scope. Universal quantifiers will feature in the attribution of both types of general belief, but in the case of

general beliefs *in sensu diviso* the universal quantifier will fall outside the scope of the belief operator, while in the case of general beliefs *in sensu composito* it will fall within the scope of the belief operator.

It seems plausible to interpret the distinction between general beliefs *in sensu composito* and general beliefs *in sensu diviso* as entailing that only general beliefs *in sensu composito* take quantified propositions as their objects. This in turn implies that general beliefs *in sensu composito* are only available to creatures that have the capacity for intentional ascent (for reasons brought out in the previous section). An analogous point holds, of course, for existential beliefs (and is, in effect, the distinction between *de re* and *de dicto* existential beliefs). A non-linguistic creature can have a belief to the effect that some F is a G in the etiolated sense that there is some F of which it believes that it is a G (with the existential quantifier falling outside the scope of the belief operator). Only a language-using creature, however, can have a belief that is correctly characterized with an existential quantifier falling within the scope of the belief operator.

Of course, the plausibility of this thesis about the language-dependence of a distinctive kind of thinking about generality stands or falls with the original distinction between two types of general belief. Some authors have put forward what appears to be an intermediate conception of general beliefs, suggesting that there is really only one type of beliefs about generality that fits into neither the distributive nor the collective mold. Braithwaite, for example, offers an account of belief according to which believing that p is a matter of behaving as if p were true. He extends this to general beliefs as follows:

A third advantage of my doctrine [the first two advantages being (a) that it lends itself to an account of degrees of belief and (b) that it can serve as the basis for an account of justified inductive belief] is that it assists in one of the most subtle problems of logic, that of the analysis of general propositions. For action appropriate to belief in a general proposition

does not present any special problems. 'I am disposed to act appropriately to every P being Q' means that, whenever I am disposed to act appropriately to a thing's being P, I am disposed to act appropriately to its being Q. (Braithwaite 1932-1933, 39-40).

Although Bennett, for example, wrongly suggests that Braithwaite is offering an account of distributive generality (Bennett 1977, 104 n.3), there is a crucial difference between Braithwaite's proposal and the distributive construal. Braithwaite does not suggest that the general belief that all Fs are G requires behaving, with respect to everything that is *in fact* an F, as if it were a G. Rather, he requires simply that the believer should behave, with respect to everything that *he takes to be* an F, as if it were a G. Failing to respond appropriately to the dangerousness of an albino tiger would not, therefore, disqualify me from believing that all tigers are dangerous (and hence the belief would not be correctly characterized with a universal quantifier outside the scope of the belief operator in the way that general beliefs *in sensu diviso* would be correctly characterized).

But it is far from clear that Braithwaite has really given us a satisfying account of general beliefs. One problem with his view is that it lacks the appropriate counterfactual implications. A believer might be disposed to behave with respect to everything that he in fact takes to be an F as if it were a G without being disposed to behave with respect to everything that he *might* take to be an F as if such things were Gs. A genuinely universal belief to the effect that all Fs are Gs is completely incompatible with the possibility that something might be taken to be an F and yet not be believed to be a G – but Braithwaite's account seems to leave this possibility open. Of course, though, a simple modification will rescue the theory. All we need require is that the believer be disposed to behave, with respect to everything that he is disposed to believe to be an F, as if it were a G. But this brings us to the real problem. Braithwaite is suggesting, not only that creatures can have fully general beliefs even though they do not have the resources to mark that generality in any way (either at the linguistic level or at the level of thought), but,

more strongly, that there are no general beliefs for which markers of generality are required. And it is this second claim that must be challenged. It cannot be the case that there is no significant cognitive difference between a creature who has a generalized Braithwaitean disposition, on the one hand, and a creature for whom such a disposition is grounded in an explicitly formulated general thought of the type expressible by a universally quantified sentence. The difference is easily capturable when we consider the logical form of the appropriate belief ascriptions. We have already seen that general beliefs *in sensu diviso* should be characterized in the following manner:

- (1) For all x , if x is an F then β believes (is disposed to act as if) x is a G

General beliefs in the Braithwaite style are specifiable in a broadly similar manner, with a slightly expanded antecedent.

- (2) For all x , if x is an F and is taken by β to be an F , then β believes (is disposed to act as if) x is a G

But this still falls short of general beliefs *in sensu composito* specifiable as follows:

- (3) β believes that, for all x , if x is an F then x is a G .

Nothing could persuade me that there is no significant difference between (2) and (3).

It would seem, then, that there is a real and genuine distinction to be drawn between two types of thinking about generality. Although one of these ways of thinking about generality (that which involves contents containing quantified general beliefs) is restricted to creatures capable of intentional ascent (and hence of semantic ascent), the distinction offers a type of thinking about generality that is indeed available at the non-linguistic level. As such, it provides a primitive analogue for quantificational thinkin

in the same manner that the forms of proto-inference discussed in Chapter Seven provide primitive analogues of the truth-functional propositional connectives.

9.6 *From domain-specificity to domain-generality*

The hypothesis of domain-specificity has become increasingly popular among theorists in cognitive archeology, evolutionary psychology and developmental psychology. (Barkow, Cosmides and Tooby 1992, Hirschfeld and Gelman 1994, Mithen 1996. But see Fodor 2000 for criticism of extreme versions of the hypothesis). This section suggests that the apparent domain-specificity of various types of non-linguistic thinking may in fact reflect a structural limitation upon the types of thinking available at the non-linguistic level.

The basic tenet of the hypothesis of domain-specificity is that certain fundamental types of cognitive activity are carried out by modular systems that have evolved to deal with particular types of problem and particular types of situation. Popular candidates for domain-specific modules include the interpersonal competences involved in social interactions; the basic principles about objects and their interactions that are usually collectively labeled naïve physics; and an intuitive grasp of folk biology and natural history. These modules operate on a highly selective and domain-specific set of inputs with a fixed and limited amount of background information. For present purposes there is systematicity within each module, but not across modules. In Chapter Six we considered an example of this failure of systematicity in the archeological record – the failure early hominids to integrate their practical abilities in tool construction with their detailed knowledge of natural history in order to produce handaxes for specific purposes. . In

the Middle Paleolithic, for example, we find what seem to be highly developed tool-making skills existing side by side with a subtle and advanced knowledge of the natural environment, but it is not until the Upper Paleolithic that we see these two bodies of knowledge being integrated in the form of tools specially designed for dealing with different plants and animal, together with hunting strategies that are tailored to the habits of specific animals (Mithen 1996a).

The proposal I want to consider is that language is required for the integration of domain-specific modules.¹⁵ The argument here hinges on the distinctive role of the relative pronoun in permitting the construction of relative clauses. This was briefly discussed in Chapter Three when we considered Quine's account of the psychogenesis of reification. I rejected Quine's argument that reification depends upon the linguistic mechanisms of quantification and the relative pronoun (the natural language equivalent of the bound variable of quantification). But nonetheless we saw no reason to dispute Quine's general claim about the enormous cognitive significance of the linguistic device of the relative pronoun. The crucial feature of the relative pronoun in this respect is that it permits the formation of relative clauses. One way of thinking about relative clauses is as a way of distinguishing within a sentence between the object that is the logical subject of the sentence (what the sentence is about) from what the sentence says about that object (Quine 1974 §24, 1995).¹⁶ In English, for example, from a sentence like 'the red deer comes to the water just before nightfall' we can extract the relative clause 'that comes to the water just before nightfall

¹⁵ A similar proposal is made in Carruthers forthcoming, although on very different grounds.

¹⁶ Geach has pointed out (1962 §§71-73) that this way of thinking about relative clauses does not extend easily to certain tricky cases. I am inclined to respond to this as Quine does (1974 p.91) by distinguishing between the most plausible account of modern English and the most plausible account of the evolution of language.

which can be used to characterize other animals, or be embedded in further sentences, and so forth. Once the relative clause has thus been constructed and detached from the original sentence the information it provides is available in a form that can be employed in a variety of different contexts. Consider a simplified model of domain-specific cognition in which a thinker can think about two separate domains, each of which contains two objects and four properties. The properties from one domain can only be applied in thought to the objects in that domain and no object or predicate features in more than one domain. Suppose that a is an object in the first domain and G a predicate in the second domain. How might such a creature come to be capable of the thought that Ga ? Only by forming a conception of what it is for something to be a G that is no longer tied to the objects in the second domain. That is to say, only by forming a conception of what it is for an arbitrary object to be a G . And it is this that the introduction of the relative pronoun makes possible. The relative clause seems the most basic way of extrapolating complex predicates from complete sentences in a form that will allow them to be applied to objects falling under other cognitive domains – and the defining feature of domain-general cognition is that objects from one domain can be thought about in terms formerly associated only with objects from another domain, as in totemic art when an artifact is attributed the properties of an animal, or as in advanced tool construction when the design of a tool is specifically tailored to properties of the intended prey.

Let us assume, therefore, that the integration of different types of domain-specific cognition must involve some cognitive mechanism that operates in a manner analogous to the relative pronoun to distinguish within thoughts between what the thought is about and what the thought affirms of that object. At this point the close relation between the relative pronoun and the variable of quantification becomes relevant once again. One might expect this cognitive mechanism to operate in a manner very similar to that involved in extracting an existential generalization from a given sentence – namely, by breaking the thought down in such a way that the predicative component is detached from the nominative component. In the case of quantified thoughts the result of this “decomposition” of the thought is to allow the

insertion of a variable bound by a quantifier into the place occupied by the name. In the case of the transition from domain-specific to domain-general thinking, however, the result of the “decomposition” is to make it possible to insert an arbitrary name into the place of the name. In both cases, however, the reliance upon the capacity for intentional ascent is clear. What is required is the ability to hold a thought in mind in order to identify and manipulate its structure. And this, by the argument of Chapter Eight, is only possible when that thought has a linguistic vehicle.

* * *

In the first seven chapters of the book I defended at some length the thesis that it can be appropriate and correct to attribute to non-linguistic creatures thoughts that have determinate contents, are compositionally structured and reflect the mode of presentation under which the creature in question apprehends the immediate environment. Not only is the ascription of such thoughts often mandated by the requirements of psychological explanation, but the explanatory practices within which such thought-ascriptions take place can be embedded within viable conceptions of non-linguistic reasoning and non-linguistic rationality.

The principal claim of the two final chapters has been that there are significant limitations to the cognitive abilities of non-linguistic creatures. Certain types of thinking are in principle only available to creatures that dispose of a language. The fundamental reason for the restriction, developed in Chapter Eight, is that thoughts can only be the objects of further thoughts when they have linguistic vehicles.

Intentional ascent (that is to say, thinking about thoughts) requires the possibility of semantic ascent. The argument from intentional ascent to semantic ascent places significant restrictions upon the scope of thinking without words. Many types of thinking involve intentional ascent explicitly, because they are directly targeted on first-order thoughts. This is the case, for example, in the types of reflection associated with second-order cognitive dynamics (9.1) and most forms of higher-order desire (9.2). Explicit intentional ascent is also involved in all forms of psychological understanding that involve attributing thoughts to others (9.3). In 9.4 I showed that intentional ascent can also be involved implicitly even in types of thinking that are not explicitly targeted on first-order thoughts – that is to say, in types of thinking whose contents do not involve further thoughts. These are types of thinking into which a creature cannot enter without the capacity for intentional ascent. The principal example discussed in section 9.4 was logical thinking involving the truth-functional propositional connectives and quantifiers (although it was also suggested that modal and tensed thought might fall into this category). In section 9.5 I identified a type of general belief that does not presuppose intentional ascent and hence is available to non-linguistic creatures. The final section of the chapter explored the role of language in making possible the transition from domain-specific to domain-general cognition. It was argued that the mechanisms required for domain-general thinking are very similar to those that make possible quantificational thought.

The picture that has emerged of the scope of non-linguistic thinking provides principled grounds for thinking that many cognitive abilities that have traditionally been taken to be uniquely human are indeed unique to language-using humans. Only language-using creatures can be logical thinkers, monitor their own processes of belief formation and argument and reflect on the desires that they want to have. Only language-using creatures are capable of attributing thoughts to other creatures. Yet the gulf between linguistic and non-linguistic thought should not be exaggerated. Many of these uniquely human cognitive abilities have analogues at the non-linguistic level. Logical thinking may be the preserve of language-users, but there are forms of general belief and types of proto-inference available at the non-linguistic

level. Non-linguistic creatures cannot monitor their own processes of belief formation, but they are nonetheless capable of sophisticated forms of belief revision. Even though the attribution of thoughts is not possible at the non-linguistic level, there are still relatively complicated ways in which non-linguistic creatures can think about the perceptions and desires of other creatures and hence in which they can explain and predict behavior in broadly speaking psychological terms. The cognitive separation between creatures that have language and creatures that do not is very real. But it is a separation between two types of *thinking* – between two ways of representing the social and physical environment – rather than between thought and the absence of thought.