SENSE, REFERENCE AND SELECTIVE ATTENTION

John Campbell and M.G.F. Martin

I—John Campbell

The idea that there is a distinction between propositional and imagistic content is familiar and compelling, but it brings with it a problem. The problem is to explain the relation between the two types of content. This problem can seem so difficult that to escape it, you would abandon the propositional/imagistic distinction. But that is an extreme reaction; we can solve the problem. My proposal in this paper is that the primary mechanism for mediating between propositional and imagistic content is perceptual attention.

Propositional content involves reference to objects. There are many ways in which we can refer to concrete objects, but the most basic sort of reference is when you can see the thing, or perceive it somehow, and refer to it on the strength of that perception. If you and I are looking out of the window, then we may discuss the castle before us, identifying it as 'that castle', the one we can see. But just having the castle in your field of view does not seem to be enough for you to refer to it. If you are to refer to the castle, you must do more than have it your field of view: you must attend to the thing. And if you are to talk to me about that castle, you have to draw my attention to it, so I get some clue as to what you are talking about. Reference on the basis of perception seems to depend on the ability to attend to the things perceived. Reference and attention are related phenomena.

Attention, as I am conceiving of it, is a matter of selection: you select some aspects of your perception rather than others. One basis for selection is perceived location: you may select the phenomena perceived at a single location. Selection on the basis of location is what I shall be calling spatial attention. Perceptual reference can depend on spatial attention. The best way for you to let me know what you are talking about may be to point to where the castle is. But there are other bases for attention, as when you and I listen to

and comment on a pneumatic drill we can hear being used, though we may have little idea where it is. I will try to determine in what sense, if any, spatial attention might be thought to be fundamental among the varieties of attention. I will also try to set out how the type of attention used bears on the sense of a perceptual demonstrative.

I begin with some remarks on the distinction between propositional and imagistic content, and the general thesis that selective attention is the primary mechanism for mediating between the two types of content.

I

Propositional vs. Imagistic Content. Propositional content is content with subject-predicate structure, in which general terms are coupled with singular terms. It is also content which stands in deductive inferential relations to other propositional contents. Not all representation is propositional. The most familiar alternative is pictorial representation. One type of pictorial representation represents the spatial relations among various objects and their parts by means of the spatial relations among the parts of the picture itself. Alternatively, the space of the picture need not be a physical space, but an array which constitutes a functional space: the array is operated on in such a way that its various components function as spatially related elements. As Kosslyn (1994) puts it, 'each part of an object is represented by a pattern of points, and the spatial relations among the patterns in the functional space correspond to the spatial relations among the parts themselves' (p. 5), and parts of the representation correspond to parts of objects. Here there is no subject-predicate structure, and no deductive inference: if you and I are deciding how to arrange the chairs in a room, our perceptions may be subjected to various transformations as we reason about what to do, but these imagistic transformations are quite different to deductive inference. On the face of it, the most basic content of perception and mental imagery is imagistic rather than propositional.

Even though propositional and imagistic representations are so different, we need some account of the relation between them. The most basic observational propositions, such as 'that cat is brown', need to be linked to perception if they are to be understood. Without

some linkage to perception, you would have no idea how to go about verifying a proposition of this sort; there would be no saying which imagistic representations serve to confirm or refute the proposition. Again, without some linkage to imagistic representations, you would have no idea how to go about acting on the basis of a proposition you believe to be true: it is because you have the imagistic representations that you know just how to move, how far and in what direction, and so on, in order to achieve some propositionally specified goal.

My proposal is that selective attention is the notion we need to describe the link between propositional and imagistic content. The general notion of attention that I have in mind is 'selection of information for further processing'. In particular, there is the selection of a body of perceptual information as all relating to a single object. In order to use an imagistic representation to verify a proposition, just the right information must be selected from the imagistic array. In order to use an imagistic representation to act on the basis of a proposition, to set the parameters for one's action, again just the right information must be selected from the imagistic array. So an account of the link between propositional and non-propositional content should focus on selective attention.

It is not, though, that we have here two independent levels of content, with an attentional link between them. The attentional link is partly constitutive of the propositional level, the attentional link is part of what makes the propositions the propositions they are. Demonstrative identification of a particular object as 'that car', for example, requires that you be selecting information from that car to use in verifying propositions about it, or in acting on the basis of propositions about it; that is what makes it the case that you are identifying that car rather than anything else. So the propositional level is partly constituted by its attentional link to the imagistic level. A parallel point holds for our grasp of general concepts. What makes it the case that you are grasping a colour concept, for example, is that you are able to select perceptual information about the colours of perceived objects in verifying or acting upon propositions involving the concept. Its having just that link to selective attention is part of what makes the concept the concept it is.

There is an asymmetry in the way in which singular terms and general concepts are linked to the imagistic level. To understand a

singular term, it seems that you must actually have selected just the right perceptual information to verify or act upon propositions involving the term. It is not enough simply to have the capacity to do that; until you make the selection of perceptual information, you do not know which thing is in question. To understand a general concept, on the other hand, it seems to be enough that you have the capacity for selective attention to that aspect of the object; it is not required that you actually should be attending to it. If you say 'that car is red', to understand you I have to select information from just the right object, but I can know what you have said without yet attending specifically to the colour of the car, though I must have that general capacity. It is a further question whether grasp of the general concept requires the ability to attend to that aspect of the thing at will, or whether it is enough if one can have one's attention drawn involuntarily to that aspect of the thing. A functioning use of propositional thought would seem to demand some general capacity for the voluntary direction of attention, but perhaps that is not needed in every single case, for every single observational concept you understand.

II

The Sense of a Perceptual Demonstrative. So much for a statement of the approach, using selective attention to characterise the link between propositional and imagistic content; can we make it do any work for us in giving an analysis of propositional content? Among the concepts used in propositional thought, I will focus on the singular terms that we use to refer to perceived objects, such as 'that car' or 'that man'. And I will consider only relatively pure uses of those terms, when they are used on the basis of current perception to refer to an object about which you have no specific prior knowledge. I said that propositional content was defined partly by the holding of deductive inferential relations between propositional contents, and a description of propositional content should give some analysis of when particular deductive relations do or do not hold between propositions. For example, the following inference seems valid: 'The Morning Star is F; the Morning Star is G; hence, the Morning Star is both F and G'. However, the following inference does not seem to be valid, though we have merely substituted for one occurrence of 'the Morning Star',

another term which refers to the same thing: 'The Morning Star is F; the Evening Star is G; hence, the Evening Star is both F and G'. This inference relies on a suppressed premise, that the Morning Star is the Evening Star, whereas the first inference needs no such premise. So there is a difference between 'the Morning Star' and 'the Evening Star'. We have the same phenomenon with perceptual demonstratives. Suppose that a single tree is visible through two separated windows of a room; some of its branches obscure one window and other branches obscure the other window. Then if I argue: 'That tree [pointing through the left window] is F; that tree [pointing through the left window again] is G; hence, that tree [pointing through the left window] is both F and G', the inference seems unproblematically valid. However, if we substitute another demonstrative which refers to the same thing, we get this: 'That tree [pointing through the left window] is F; that tree [pointing through the right window] is G; hence, that tree [pointing through the left window] is both F and G'. And this inference is invalid; it has to be supplemented by an extra premise saying that it is one and the same tree both times. So there is more to the meaning of a perceptual demonstrative like 'that tree' that just that it refers to the object it does. We ought to be able to say when this kind of trading on identity is legitimate and when it is not. This is the problem of sense. Sense is that, sameness of which makes trading on identity legitimate, difference in which means trading on identity is not legitimate. So far as I know, the problem of the senses of perceptual demonstratives was first raised by David Kaplan in 'Demonstratives'. He said:

> for a Fregean the paradigm of a meaningful expression is the definite description, which picks out or denotes an individual, a unique individual, satisfying a condition *s*. The individual is called the *denotation* of the definite description and the condition *s* we may identify with the *sense* of the definite description. Since a given individual may uniquely satisfy several distinct conditions, definite descriptions with the same sense may have the same denotation. And since some conditions may be uniquely satisfied by no individual, a definite description may have a sense but no denotation. The condition by means of which a definite description picks out its denotation is *the manner of presentation* of the denotation by the definite description.

> The Fregean theory of demonstratives claims, correctly I believe, that the analogy between descriptions (short for 'definite descriptions') and demonstrations is close enough to provide a sense and a denotation analysis of the 'meaning' of a demonstration. The denotation is the demonstratum (that which is demonstrated), and it seems quite natural to regard each demonstration as presenting its demonstratum in a particular manner, which we may regard as the sense of the demonstration. The same individual could be demonstrated by demonstrations so different in manner of presentation that it would be informative to a competent auditor-observer to be told that the demonstrata were one. For example, it might be informative for me to tell you that
>
> > That [pointing to Venus in the morning sky] is identical with that [pointing to Venus in the evening sky].
>
> (I would, of course, have to speak very slowly.)
>
> (Kaplan 1989, p. 514)

The problem this raises is how we are to characterise in detail the senses of perceptual demonstratives, such as 'that planet', or 'that car'. In the case of a definite description the phrase itself makes explicit the condition that something has to meet for an object to be its denotation. If two descriptions impose just the same conditions, then trading on identity is legitimate. But how are we to say what the sense of a demonstrative is?

This problem of the sense of a perceptual demonstrative is a problem about selective attention. To find when two demonstratives have the same sense, we have to look at the principles that the perceptual system uses to select a collection of imagistic information as all relating to a single object. The use of a demonstrative depends on some principle of selection being used to isolate some of one's current imagistic information as all relating to one object. When we have two demonstratives that depend on the same imagistic information having been selected using just the same principle, then we have sameness of sense and the identity statement involving those demonstratives will be uninformative. If the imagistic information has been selected on different basis, then the senses will be different and the identity will be informative. Another way to put my proposal is in terms of the idea that visual processing involves the use of 'feature maps': that the various features of the objects one perceives, such as colour, shape or movement, are processed separately by the visual system, which

then has the problem of binding together the features which are features of the same object. Anne Treisman (1993) has proposed that the binding is achieved in selective attention. It follows from this that there is as yet no binding of features in the unattended areas of an imagistic representation. So demonstrative identification of an object requires the exercise of attention, using some principle or principles to bind together features of the same object.

What are the principles that we use to select a body of information as all relating to a single thing? One fundamental method of selection that we use is location: each of the pieces of perceptual information we have about various features such as colour, shape and so on has a location implicitly assigned, and information designated as coming from the same location is selected as all relating to a single object. So Treisman (1993) argues that in vision, the binding together of information from different feature maps is achieved by the system making explicit the common location of all the features implicitly designated as at a particular place. When you try to attend visually to an item on the basis of some other feature that it has, such as its colour or its size, your selection of that item has to be mediated by its location. That is, you may be trying to spot the red one, or the biggest one, but you can follow through on that only when you see where the thing is, and only by seeing where the thing is (Nissen 1985).

The implication of this is that demonstratives which depend on spatial attention—attention which uses location as the principle of selection—will have their senses individuated by the locations used in selecting the underlying collections of information. We can use this point to explain the distinction between cases in which trading on identity is legitimate and cases in which trading on identity is not legitimate. The question is whether the same principle is being used to bundle together the underlying perceptual information as all true of a single object. So in the case of demonstratives referring to a tree, which rely on spatial attention, the question is whether the demonstratives are using the same location to bundle together the information as relating to a single thing. Pointing through the same window indicates that the same location is being used; pointing through different windows indicates that different locations are being used. The notion of location used here is itself relative to a frame of reference:

sameness or difference of location means sameness or difference in the relevant frame.

I think it is instructive to contrast what I have said so far with Evans' (1982) account of demonstrative reference. The key to Evans' account is his view that possessing a concept of the property of being F is knowing what it is for a proposition of the form '∂ is F' to be true, where '∂' is a fundamental idea of an object. A fundamental idea of an object identifies it as an object of a certain sort: as, for example, a number, a shape, or a mountain. It further identifies it as the possessor of the characteristics which ultimately differentiate it from all other things of the same sort. So, a number is ultimately differentiated from all other numbers by its position in the number series; a spatial object is ultimately differentiated from all other things of the same sort by its location at a time, its position then with respect to other objects. For our purposes, the key point here is that this way of identifying an object is not demonstrative; it is meant to locate the thing 'objectively' (Evans speaks of identification at the level of a cognitive map). This raises the question how there can be such a thing as non-fundamental singular reference. For if predicative concepts are explained and introduced at the fundamental level, how could your grasp of a predicative concept F and your grasp of a non-fundamental idea *a* combine so that you can grasp the thought that *a* is F? Evans' answer is that what constitutes your understanding of a non-fundamental idea *a* is your knowledge of what would make true an identity of the form '*a* is identical to ∂', where '∂' is an arbitrary fundamental idea. So an understanding of a perceptual demonstrative, such as 'that man', will consist in knowledge of what would make true an identity of the form 'that man is identical to ∂', where '∂' is an identification of the thing by its location; and it is perception of the egocentric location of the thing, according to Evans, that provides one with that knowledge. But this analysis of why location is central seems quite mistaken. It is very hard to make sense of the 'fundamental level of thought' about spatial objects Evans describes, however plausible it might be to suppose that there is such a canonical level of thought about abstract objects such as the natural numbers. And even if there were such a level of thought, the idea that predicates applicable to concrete objects are first introduced and explained at the fundamental level is

indefensible: observational predicates are first introduced and explained in the context of perceptual-demonstrative thoughts. The insight in Evans' account is that perception of location is central to grasp of the perceptual demonstrative. But the kind of *a priori* argument he gives seems bound to fail. The centrality of perceived location for the sense of a demonstrative is rather, I have suggested, a consequence of the empirical fact that in vision at any rate, perceived location is of pervasive importance for selective attention. I will pursue this point in §III.

So far I have emphasised the role of spatial location in demonstrative reference. This can only be part of an account of demonstrative reference, because so far I have not taken any account of the causal structure of physical objects. It seems that in order to be referring to physical objects as physical objects, you must have some understanding of the possibility of causal interaction between them. Someone who sees only variously coloured patches of light, and who has no conception of the possibility of interaction between them, is not in a position to make reference to them as physical objects. The mere capacity for spatial attention is not enough for this grasp of the causal significance of objects; you can use location to single out features at a location without yet having any understanding of their causal significance. One way to make the point is to remember that there can be distinct objects in the same place at the same time. For example, a hammer, a mass of molecules and a particular block of wood and metal are all in the same place at the same time. If you are to refer to one rather than the other, you need to use more than spatial location; spatial location alone will not differentiate between them. If you ask what does make the difference, it seems to be causal structure: the causal principle of unity of the thing, the way in which the condition of the thing at one time depends on its condition at earlier times. So spatial attention is not enough for singular reference to an object; we need some grasp of its causal structure. I think that by looking at the way in which we grasp causal structure, we can achieve some understanding of why propositional thought has the architecture that it does.

The most basic way in which you can display some appreciation of the causal structure of an object is by the way you act on it. There can be a lot of structure in this. Of course, there are relatively simple

actions, such as eating the thing or throwing it away, which need not involve much in the way of an understanding of its causal structure. But there are also what I will call manipulations of the object, which require that you have some grasp of which properties the thing has, and which involve you in changing one of the properties of the thing in order to affect its further characteristics. So, for example, you might squash something to get it to fit, rub two sticks to start a fire, or oil a wheel so that it runs freely. In these cases you are using your grasp of the properties of the object and their causal significance: you know which properties the thing currently has and you know how modifying these properties will affect the behaviour of the thing. There is also the way in which your perceptions of the object can be operated on in imagistic reasoning. This too may exploit your knowledge of the properties of the thing and your grasp of their causal significance. But the significance of this reasoning is ultimately in its implications for your actions.

We can call this grasp of the causal structure of the object a practical grasp; and we can contrast it with the theoretical grasp of causation that we have at the level of propositional thought. This consists in the ability to give causal explanations of the phenomena we observe. At the level of a practical grasp of causation, there is no need for any capacity to give causal explanations. But at the propositional level, there is a capacity to say what the properties of the object are, and to use them in giving explicit causal explanations. Giving explicit causal explanations requires the use of something like propositional content: we need subject-predicate structure and the possibility of deductive reasoning. We need the ability to refer to individual objects, to ascribe predicates to them and to ascribe the same predicates again and again, so that a grasp of physical law can be put to work. If you are constructing a robot, then the kind of computational system you select for it to use will depend on what kinds of computations you want it to perform; you can explain how the particular computational system will be exploited by the robot. Similarly, we can ask why humans should use a representational system, with just the features of subject-predicate structure and use of deductive inference, that I began by describing. What is the fundamental type of human reasoning that exploits this structure? I am suggesting that it is precisely this type

of reflective causal thinking that demands and exploits subject-predicate structure and deductive inference.

If we think of attention as selection for further processing, then these points relate to the type of further processing that is in question. There are two types of further processing that display grasp of the causal structure of physical objects: action on those objects, by manipulating them in various ways; and propositional reasoning in which one gives causal explanations for the behaviour of those objects. Of course, there must be some harmony between those two aspects of your grasp of causal structure, and it must be possible for them to affect one another: what you figure out in giving causal explanations must be capable of affecting your actions with respect to the object, and your trial and error actions on the object must be capable of affecting the causal explanations you give.

III

Auditory Demonstratives. It is easy to feel that there is some inevitability about the role of spatial attention in demonstrative reference, given that the reference is to spatiotemporal objects. I think that this is a mistake, which comes in part from acceptance of inconclusive *a priori* arguments that things must be so, and in part from an exclusive concentration on visual demonstratives; for location is certainly, as a matter of empirical fact, fundamental to visual attention. So far, I have been talking principally about visual demonstratives; I think it is helpful now to look briefly at auditory demonstratives. There is no immediately obvious reason to suppose that spatial attention and object reference work in the same way in all the sensory modalities. We can begin by making a broad division between two types of auditory demonstrative. Suppose that outside in the street you hear a bulldog and a pekingese start to fight; you hear the whole thing from the initial growls and yaps to the final triumphant or defiant yowls, and can tell exactly when the one dog leaves off and the other begins. Well, you might say, that bulldog put up a good fight, but it had no chance. Here the demonstrative ‘that bulldog’ is an auditory demonstrative—you may at no point have looked at the scene—but it refers to an ordinary physical object. In contrast, if you simply hear a grinding noise in the street outside and have no idea what the source is, you

may simply refer to the noise itself, you may say that you wish that noise would stop. In both cases, selective attention seems to be the mechanism of reference. In the case of the bulldog, it is what ties together all that noise as emanating from a single animal; in the case of the grinding, it is what lets you single out that noise rather than any other. But is there anything particularly spatial about the exercise of attention in these cases?

The problem you have, in constructing the auditory demonstrative, 'that bulldog', is to put together just those bits of auditory information which come from a single source, the dog in question. When solving the parallel problem for vision, it obviously is an excellent strategy, given the kind of environment we occupy, to put together information from the same location. But for sound the same strategy is not likely to be so useful. It is true that we can use the differences in the onset time of a tone at the two ears to determine a direction for the sound, but even in favourable conditions auditory localisation is not particularly precise. But anyway auditory localisation is not a reliable guide to the location of the source of a sound. To determine the difference in the onset time of a tone at the two ears we have to know when it is the same tone at each ear, and outside the laboratory, in a noisy environment, that is not easy. And anyhow, sounds can be echoed or bounced off walls, can be deflected by intervening obstacles, so that the location determined by this strategy may not be where the source of the sound is. And these factors may operate differentially on different sounds produced by the source, so that spatial bundling is not an effective way of putting together all the auditory information that one has from a single source (Bregman 1993).

Much of auditory attention may be schema-based. For example, if your dog has been in many fights, you may know exactly how he sounds in one, whereas the lay person listening to the scene might have some trouble in sorting out where one dog leaves off and the other begins. You have built up a schema for your dog's performance, which you can use to filter out all but the auditory information coming from it, a schema which the lay person lacks (Jones and Yee 1993). But it also seems that there are general principles which can be exploited to group sounds as deriving from the same source, even the absence of a specific schema. For a simple example, suppose you are presented with a steady sound

which suddenly becomes more intense for a brief period and then suddenly returns to its original intensity. This is typically heard as involving the initial sound being joined by another for a brief period, the original sound continuing at its original intensity throughout. The auditory system discerns sounds from two sources here. On the other hand, if the onset of the increase in intensity is gradual rather than sudden, this is heard as the original sound having changed in loudness, so that only one source is discerned. This displays some simple non-spatial principles of auditory organisation at work. And these principles can affect the perceived location of the sound. Suppose we modify the example just given, so that the right ear is presented with a tone just like the one already described. Simultaneously, the left ear is also presented with a sound just like the one already described, except that it always stays at the lower intensity. So initially the subject hears a single tone coming from the middle. When the tone at the right ear suddenly increases in intensity, the subject hears the original tone as continuing to come from the middle, but being joined briefly by another tone at the extreme right; so he hears two tones coming from different places. In contrast, suppose the intensity of the tone at the right ear is raised and lowered gradually. Then the subject hears only a single tone throughout, but one which gets louder and moves over to the right before swinging back to the middle. This example, from Bregman 1993, shows that more fundamental principles of auditory grouping may determine the perceived location of the source of a sound; perceived location is not a principle more fundamental than any other in putting together various pieces of auditory information as deriving from a single object.

I have said that the importance of location for visual demonstratives is a consequence of the empirical fact that location is fundamental to visual attention, and that if we consider auditory demonstratives such as ‘that bulldog’ we find that location need not be fundamental to demonstrative reference. You might wonder whether there is not, nonetheless, some *a priori* argument to be given for the importance of spatial attention for reference. Let me separate two ways in which you might try to argue this point. One is to argue that there is something *a priori* fundamental about spatial attention among the forms of attention; that we can make

nothing of the idea of people who perceived a world just like ours and who engaged in demonstrative reference very much as we do, but who lacked the capacity for spatial attention. I shall consider this line of thought in more detail in a moment. But immediately, we can see that it would be possible to acknowledge that there is something *a priori* fundamental about spatial attention without acknowledging that demonstrative reference must therefore be based on spatial attention; everything depends on just why spatial attention is held to be *a priori* fundamental. A second line of argument is to acknowledge that auditory demonstratives characteristically do not depend on spatial attention, but to maintain that auditory reference to concrete objects is somehow parasitic upon visual or tactual reference to concrete objects. Straight off we can see that the problem with this second line of thought is that you could accept that auditory reference is parasitic on visual or tactual reference without thinking that this has anything to do with the types of attention used by the perceptual systems. Hearing would depend on vision or touch even if auditory attention were exclusively spatial attention. For example, it might be argued the key point is that vision and touch tell us about properties of the objects perceived which are causally more basic than those we find out about through hearing. Listening to an object can tell you a lot about it. The mechanic listening to the car engine ticking over, or the doctor listening to his patient cough, can thereby find out a lot about the behaviour of the thing. But it is arguable that these properties are in a sense less fundamental to the objects in question that the basic volumetric properties—size, shape, solidity and so on—about which vision and touch inform us; it might be said that it is only through vision and touch that we have information about the properties of the object in virtue of which it counts as a space-occupant at all (Evans 1984). But even if this line of thought is correct, it evidently does nothing to show that the reason why vision is more fundamental than hearing has anything to do with the spatial character of visual attention.

What about the view that there is, *a priori*, something fundamental about spatial attention among the forms of attention? This does seem a compelling idea, but it is not easy to explain just why. One proposal that might be made is that spatial attention is necessary for the perception of spatial relations between objects

(Logan 1994). Another argument that might be given is that spatial attention is used by the most basic ways we have of drawing one another's attention to this or that object. If we are to have voluntarily directed attention to objects, we need some way of specifying the target of the attention without presuming the object to be in view already, and that is precisely what spatial attention provides (Logan 1995).

To sum up. Spatial attention may be particularly important in communication involving perceptual demonstratives. And there may be a case to be made for saying that spatial attention is fundamental among the forms of selective attention, for spatial attention is needed to determine the spatial relations between perceived items. And an ability to perceive the spatial relations between perceived items might be thought to be needed for reference to spatiotemporal objects. It remains an empirical matter which types of perceptual demonstrative depend on spatial attention. There can be no *a priori* argument that somehow spatial attention is the only type of attention which can ground demonstrative reference. There are many types of selective attention which can mediate between propositional and imagistic content.

I end with some remarks on the interpretation of this approach to demonstratives.

IV

Immunity to Error Through Misidentification. Any account of demonstratives should address the phenomenon of immunity to error through misidentification. For definiteness, I will discuss this in relation to demonstratives based on the use of spatial attention. Suppose we are listening to sound through a stereo system and one of the speakers fails. I point to the left-hand speaker and say 'that speaker has gone dead', but in fact it was the other speaker that had gone dead. Here I was right about the predicative part: one of the speakers had indeed gone dead. My mistake was a mistake of identification: I was wrong about which speaker it was. So this kind of case contrasts with the case in which I make a mistake of predication, when both speakers are working normally but something has gone wrong with my hearing. Incidentally, in the case here in which I make a mistake of identification, this does not

depend on there having been some inference underlying my judgement; in this case there is no inference underlying the judgement. There do seem, though, to be cases in which error through misidentification seems impossible. For example, if I judge, 'that speaker is just in front of me, slightly to the right', it cannot ordinarily be that I say that because some speaker is just in front of me, slightly to the right, and I have just made a mistake about which one it is. An account of the sense of perceptual demonstratives should explain why some judgements about perceived objects are subject to error through misidentification and some are not.

Just to repeat the phenomenon: if you judge, 'that chair is yellow', it may be that you thereby know of something that it is yellow, but that thing is not the chair, if, for instance, the chair is transparent and set against a yellow background. If your judgement is mistaken, you can rectify that by retreating to the more cautious, 'At any rate, something is yellow'. You made a mistake of identification. But there are cases in which a mistake of that sort is impossible. If I judge that the chair is right there, in front of me and to the right, it can't be that I make that judgement because something is in front of me and to the right, only that thing is not the chair. If I have made a mistake about location, I can't rectify the situation by saying, 'At any rate, something was there'. The distinction here is between the characteristics of the object, such as its location, which you use to select it visually, and which determine which thing you are talking about, and characteristics such as colour, which do not play such a fundamental role in selection.

What makes an error of identification possible is that you are selecting imagistic information as all relating to a single object, and you might make a mistake in doing that, in that you bundle together information which actually relates to different objects. But the specification of location which is the basis for your selection of the information as all true of a single object cannot relate to some other object than the one you identify. It is the principle which ties all the information together as true of a single thing, so there is no way in which it could somehow be assigned to another bundle of information, which in its turn would all be bound together as relating to a single location.

I should not overstress the analogy between demonstrative senses, on the view of them as attention-based, and the senses of definite descriptions. I introduced the notion of sense in connection with trading on identity in deductive inference. Sense, on this account, must determine reference, in that sameness of sense implies sameness of reference. For otherwise, we could have two terms with the same senses but different references, in which case it would not be legitimate to trade on identity in an inference using the two terms; but sense was just introduced as that sameness of which means trading on identity is legitimate. So in the case of demonstratives based on spatial attention, sameness of perceived location must imply sameness of reference. In the case of a definite description, though, a stronger reading of the thesis that sense determines reference also holds true: the sense gives a statement of what makes it the case that a particular object is the referent of the term. This does not seem to apply to demonstratives. You might propose that a demonstrative based on spatial attention must mean something like, 'the object at position X', where the position is being identified perceptually. If you descriptively identify an object by means of its location, then there is no prospect of getting it wrong about the location of that thing: a mistake about location can only mean that you have failed to identify anything at all. To say that location is the basis of selective attention, though, is not to say that you cannot get it wrong about the location of an object. You can use apparent location to select an object even though it is not where it seems to be. Suppose, for example, that you see an object in a mirror without realising that there is a mirror there. You might use the apparent location of the perceived features as the basis for selecting them as all features of a single thing, and succeed in attending to that object, and consequently, in perceptually demonstrating the object, even though it is not where it seems to be. It is still true that the location you use to select the information gives the sense of the demonstrative you use, and that any two demonstratives using the same location as the basis for selection will have the same sense, even though the object does not in fact have that location.

You might say that this shows only that the relevant descriptive condition is not 'the object at position X', but 'the object which looks as if it is at position X'. The point to notice about this is that

there is something metarepresentational about this use of 'looks'. We might gloss 'the object which looks as if it is at position X' as: 'the object which the visual system represents as being at position X'. And to determine which object that is, it is not enough to find which object actually is at position X. We have to look at the causal source of the visual representations, and ask whether that object matches the features bundled together by spatial attention sufficiently well to be the object perceived, rather than the cause of an hallucination.

The contrast between a descriptivist and an attention-based account of demonstrative senses matters when we consider the dynamics of demonstratives: our ability to keep track of objects over time, as we or they move around, and our capacity to keep track of objects across sensory modalities, as when we take the object held to be the same as the object seen. On a descriptivist account, it seems that the senses must be different, since the perceived location of the object is changing moment by moment, so different locations are being assigned to it. An attention-based account, however, can acknowledge that what we have here is a single temporally extended principle for binding together perceptual information as all true of a single object, a principle that can also extend across modalities. So we can have a single demonstrative sense in play when we are keeping track of an object across time or across modality.

One way to look at the situation is to say that in demonstrative reference at the level of propositional thought, the interest is, of course, in objects rather than in locations, so change of location is unimportant so long as it is transparently the same object that is in question. But there is something further to be said about how it can be transparently the same object that is in question, given that the visual system uses location to single out objects and that the location of the object may have changed. We could put the point here by distinguishing between attention and the control of attention (Treisman 1993). On this view, attention—the selection of information for further processing—is always spatial, is always achieved by selecting the information from a particular location for further processing. But there is also the question how we go about deciding which location to designate, as the location information from which will be selected. That is the control of attention. On this

view, attention to a particular location is what bundles together various features as features of a single object, in the first place. But then you can use that bundle of features to determine what location to attend to next—the visual system simply finds the location which has that bundle of features. And of course a range of our principles, relating to the kinds of trajectory an object may be expected to have, and the kinds of deformation that it may expected to undergo while in transit, will also come into play here.

Similar remarks may apply to keeping track of objects across sensory modalities. Here, though, there is a further question about whether there is a single frame of reference being used by vision, touch and hearing, or if these senses each use their own frame of reference with linkages in the control of attention between them (for discussion see, for example, Driver and Spence 1994, and Levinson 1996). Whichever way we go on that point, it can still be true that although the sense of a demonstrative can be individuated by a location at a time (and a sensory modality), this does not provide a descriptive sense for the demonstrative, but rather, a cross-section of the capacity to keep track of an object, across time or across sensory modality.

I have developed the view this far partly in order to show the range of issues that come into play when we approach the problem of demonstrative senses in terms of the cognitive science notion of selective attention. It seems to me that the foregoing remarks only begin on the depth of understanding we can achieve by taking demonstrative reference to be a phenomenon of attention. And, of course, as I said at the outset a parallel approach could be developed for our understanding of observational predicates: that, too, involves an ability to attend selectively to one object rather than another, but it also involves the capacity to select, from among the information you have from that object, just that information which relates to the applicability of the relevant observational predicate. The problem of understanding the relation between concepts and perception is about as fundamental as any which we currently face, and the key to it is provided by the notion of selective attention.[1]

1. For discussion of earlier drafts, I am grateful to Bill Brewer, Naomi Eilan, Alison Gopnik, Ervin Hafter, Christopher Peacocke, and John Watson. Earlier versions were presented to seminars in Oxford and to Berkeley Psychology Department, and I am indebted to participants for their comments.

REFERENCES

Bregman, Albert S. 1993. 'Auditory Scene Analysis: Hearing in Complex Environments'. In Stephen McAdams and Emmanuel Bigand (eds.), *Thinking in Sound* pp. 10–36. Oxford: Oxford University Press.

Driver, Jon and Charles J. Spence. 1994. 'Spatial Synergies between Auditory and Visual Attention.' In Carlo Umiltà and M. Moscovitch (eds.), *Attention and Performance* vol. XIV, pp. 311–331. Hillsdale, N.J.: Erlbaum.

Evans, Gareth. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.

Evans, Gareth. 1984. 'Things Without the Mind'. In his *Collected Papers*, pp. 249–290. Oxford: Oxford University Press.

Jones, Mari Riess and William Yee. 1993. 'Attending to Auditory Events: the Role of Temporal Organisation'. In Stephen McAdams and Emmanuel Bigand (eds.), *Thinking in Sound* pp. 69–112. Oxford: Oxford University Press.

Kaplan, David. 1989. 'Demonstratives'. In Joseph Almog, John Perry and Howard Wettstein (eds.), *Themes from Kaplan* pp. 481–563. Oxford: Oxford University Press.

Kosslyn, Stephen M. 1994. *Image and Brain*. Cambridge, Mass.: MIT Press.

Levinson, Stephen. 1996. 'Frames of Reference and Molyneux's Question'. In Paul Bloom, Mary A. Petersen, Lynn Nadel and Merrill F. Garrett (eds.), *Language and Space*, pp. 109–169. Cambridge, Mass.: MIT Press.

Logan, Gordon D. 1994. 'Spatial Attention and the Apprehension of Spatial Relations'. In *The Journal of Experimental Psychology: Human Perception and Performance* vol. 20, pp. 1015–1036.

Logan, Gordon D. 1995. 'Linguistic and Conceptual Control of Visual Spatial Attention'. *Cognitive Psychology* vol. 28, pp. 103–174.

Nissen, Mary Jo. 1985. 'Accessing Features and Objects: Is Location Special?' In Michael I. Posner and O.S. Marin (eds.), *Attention and Performance* vol. 11, pp. 205–219. Hillsdale, NJ: Erlbaum.

Treisman, Anne. 1993. 'The Perception of Features and Objects'. In Alan Baddeley and Lawrence Weiskrantz (eds.), *Attention: Selection, Awareness and Control*, pp. 5–35. Oxford: Oxford University Press.