CHAPTER 20

WIDE CONTENT

FRANCES EGAN

THOUGHTS are individuated both by their attitude type and by their content. A thinking subject might bear different attitudes to the same content—she might, for example, both believe and fear *that the conflict in the Middle East will not be resolved*—or the same attitude to different contents—she might believe, say, *that the Devils will win the Stanley Cup* and *that it will snow tomorrow*. *Wide* or *externalist* content is individuated in part by reference to features of a subject's external surroundings, by her physical environment or the community with which she shares a language. Externalists hold that physically identical subjects might have thoughts with different, wide, contents. *Internalists* either deny that content is wide, claiming that all content supervenes on intrinsic physical states of the subject, and is hence *narrow*, or insist, at the very least, that the content that individuates a subject's thoughts in predictions and explanations of her behaviour must supervene on her intrinsic properties.

I begin by discussing the arguments for and against wide content. I then argue that genuinely explanatory content is wide.

20.1 ARGUMENTS FOR WIDE CONTENT

The existence of wide content is established by several well-known arguments. The arguments share a common structure: we are asked to consider physically identical subjects in various contexts where, it is suggested, certain features of the context determine different contents for the subjects' thoughts.

Hilary Putnam (1975) argued that the meaning of natural-kind terms is determined, in part, by the nature of the subject's physical environment. Putnam asks us to

consider a world, Twin Earth, in which the vocable 'water' is used to refer to the liquid filling this world's oceans and streams, a liquid superficially similar to water, though chemically quite different. Call this liquid 'XYZ'. Imagine two subjects, Oscar and his physically identical counterpart, Twin Oscar, circa 1750, before the development of modern chemistry.[1] Each subject utters the form of words 'Cold water is refreshing on a hot day'. According to Putnam, the meanings of the two utterances are different, in virtue of the fact that they are used to refer to different substances—H_2O by Oscar and XYZ by Twin Oscar. If Putnam is right, the internal, psychological state of the speaker is insufficient to determine the meaning of natural-kind terms; the nature of the environment also plays a meaning-determining role. As Putnam famously put it, 'meanings just ain't in the head'. Moreover, on the assumption that a subject's utterances express his thoughts, the twins' thoughts have distinct contents, in virtue of the distinct concepts expressed by their use of the vocable 'water'.

Tyler Burge (1978) argued that the contents of many thoughts are determined, in part, by community standards of correctness. A medically unsophisticated subject, Alf, believes that he has arthritis in his ankles, wrists, and, as it happens, his thigh. When informed by his doctor that arthritis is a disease of the joints, Alf gives up the belief that he has arthritis in his thigh. We can compare Alf with a second subject, Twin Alf, who has an identical history, non-intentionally described. Twin Alf belongs to a linguistic community that uses the vocable 'arthritis' to refer to an ailment that affects the thigh as well as the joints. Both Alf's and Twin Alf's use accord with this community's use. Burge concludes that the beliefs that Alf and Twin Alf would express using the form of words 'Arthritis is painful' have different contents. (On the assumption that beliefs have their contents essentially, Alf and Twin Alf have different beliefs.) The content of thoughts ascribable using general terms such as 'arthritis' is determined, in part, by community-wide social and linguistic practice, and is therefore wide.

An argument originally due to Gareth Evans (1982) establishes that thoughts involving singular reference—so called object-involving thoughts—have wide content. My belief that *this glass* (imagine me pointing to a specific glass in front of me) is empty depends constitutively upon the existence of a particular glass. My perceptual contact with this glass puts me in a position to make a singular demonstrative judgement that *this glass is empty*, a judgement that my Twin Earth counterpart, pointing to a different glass, is not in a position to make. Her judgement has a different wide content. The content of object-involving thoughts is the least controversial type of wide content and will not be discussed further.

20.2 Arguments against Wide Content

The externalist interpretation of the thought experiments has been disputed by many. Internalists argue that the twins' thoughts share a single *narrow* content

[1] As many commentators have noted, we must ignore the fact that humans are largely composed of water.

no matter what the context. Because the Putnam and Burge arguments appeal to different aspects of the thinker's context—the physical environment and the social or linguistic context respectively—they raise somewhat different issues and so will be discussed separately.

It has been suggested that the difference in Oscar's and Twin Oscar's physical environments does not give rise to a difference in the contents of their thoughts, because both believe something like the following: *the clear potable liquid that fills the oceans and streams is refreshing*. The revised belief contents are assumed to supervene on intrinsic properties of subjects. The issue, more generally, is whether the twins' concepts include within their extension not only the liquid that fills their local oceans and streams (H_2O and XYZ respectively) but also anything phenomenally indistinguishable from this liquid. If, as Crane (1991) and Segal (2000) have argued, the twins' concepts apply to the 'watery' substance in their worlds in virtue of the relatively superficial properties shared by the two substances, then they have the same concept, and Putnam's argument fails to establish the existence of wide content.

The fact that nobody in 1750 was in a position to provide a distinguishing description of the two liquids appears to undermine the claim that the concepts in question express natural kinds, or kinds distinguished by a hidden nature, as Putnam's argument requires, rather than more superficial kinds or 'motleys', what Locke called 'nominal kinds'. This, however, is a mistake. As Putnam (1975) argued convincingly, whether the twins' terms express natural-kind concepts or more superficial kind concepts is an empirical question turning on the referential intentions of the speaker. Archimedes knew nothing about the chemical composition of the substance he called χρυσος, but if he would have rejected as a genuine instance of the kind something that possessed only the observable properties of gold—if he would have agreed that all that glitters is not gold—then χρυσος for Archimedes indeed expressed a natural-kind concept (arguably, the concept *gold*). The speaker need have only a vague conception of a shared hidden nature or structure determining the kind; she need not have any idea what that hidden nature in fact is. Whether a speaker has such a conception in mind can normally be determined by investigating the speaker's use of the term under a wide range of circumstances. This process would include querying the speaker about what she would say in various counterfactual circumstances. ('If it looked like gold but in fact was made out of a high-tech synthetic would it be gold?') We can imagine a population which, having no interest in the hidden nature or structure of objects and substances, intends by their use of such vocables as 'water', 'gold', and 'aluminium' to refer to things that possess only the observable properties of water, gold, and aluminium. Their use of these terms would be much like our use of 'table' or 'chair', expressing relatively superficial kinds (in this case, functional kinds) or motleys, rather than natural kinds.[2]

[2] These people would have little interest in basic science, which, of course, seeks to explain the observable properties of things by appeal to underlying nature or structure. We might imagine them to be served by a distinct population responsible for developing and maintaining the useful objects that keep the society running. This group would speak a different language, with genuine natural-kind terms.

It is worth noting, however, that a speaker's (or a community's) referential intentions may not be determinate enough, in advance of theory development and scientific discovery, to determine whether a particular sample falls within the extension of a given concept. The notion of hidden nature or structure that a speaker has in mind and that governs her use of a concept may be too vague to settle whether objects with *similar* microstructures (say, isotopes of the same element) would fall under the speaker's concept. But indeterminacy in the concept's extension does not, as Segal (2000) suggests, turn it into a motley concept.[3,4]

The externalist interpretation of Burge's thought experiment has also been challenged. Crane (1991) and Segal (2000) have argued that a single content ascription is true of Alf and Twin Alf, one that reflects what we might call their *personal conception* of how things are. Alf's concept, on this view, is idiosyncratic, relative to his linguistic community; it is the same concept as Twin Alf's. (Twin Alf is assumed to be a linguistically competent member of his own linguistic community.) We might call this shared, *narrow*, concept 'narthritis', to distinguish it from *arthritis*, the concept that applies only to the joints. Burge's assumption that the community's concept should be attributed to Alf despite his partial or inadequate understanding of that concept is simply wrong, on this view. Rather than having a false belief that he has arthritis in his thigh, Alf believes, truly, that he has *narthritis* in his thigh, that the *narthritis* in his ankles is getting worse, that aspirin relieves the symptoms of *narthritis*, and so on. He also has the false metalinguistic belief that the vocable 'arthritis', in the mouths of his cohort, expresses the concept *narthritis*.

As noted above, whether a speaker's use of a term expresses a natural-kind concept can normally be determined by investigating the speaker's behaviour under a wide range of circumstances. It is not clear that the present dispute can be settled in quite as straightforward a fashion. Burge cites Alf's 'post-correction' behaviour—specifically the fact that when informed by his doctor that arthritis is a disease of the joints Alf regards himself as having been mistaken about arthritis—as evidence that his concept, all along, is the same as the community's. But an internalist interpretation of Alf's post-correction behaviour is just as plausible: Alf does intend to deploy in speech and thought the same concepts as his cohort, but his deference to the doctor should be construed not as evidence that prior to correction he did in fact deploy the community's concept, but rather as a *change* in concept, bringing his own usage into line with his cohort. In any event, the situation seems quite unlike the natural-kind case. It is not clear that anything in Alf's behaviour, or dispositions to behaviour (including his disposition to retract his claim about his thigh if experts tell him that one can't have that sort of ailment in the thigh) is sufficient to determine whether, prior to correction, the belief that Alf would express by his use of the vocable 'arthritis is painful' has the same content as it would have in the mouths of his cohort, or whether it has some idiosyncratic content reflecting his personal conception of things.

[3] It is likely that the shapes of most of our natural-kind concepts are in flux, as empirical discoveries and developments in theory force a revision of their extensions.

[4] For further discussion of the issues in the last two paragraphs see Brown (1998).

It might be objected that the internalist interpretation of Alf's post-correction behaviour requires, implausibly, that a subject's concepts can change without the subject being aware that they have done so.[5] But this is not an objection that an externalist can make. The use of the vocable 'arthritis' in Alf's community might shift somewhat without Alf's knowledge. On Burge's view, Alf's concept would change, even though he isn't aware of the change. Even if the concepts in question are assumed to correspond to the subject's personal conception of things, there should be nothing troubling in the idea that a subject's concepts can change without her being aware of that fact. 'Concept' is a theorist's term of art. Alf is fully aware that his *beliefs* have changed. After correction by the doctor he is disposed to utter the form of words 'I no longer believe that arthritis can occur in the thigh'. On some accounts of concepts—in particular on the internalist account under consideration, but more generally on holist conceptions of concepts—this change in Alf is sufficient to constitute a change in his concept, but there is no reason to think that the subject should be aware of *that*.[6]

Some authors have argued that the thought-experiments do not force a choice between wide and narrow content; in some circumstances it is correct to ascribe both kinds of content, each answering to different practical and theoretical interests. According to Loar (1988), Alf and his linguistic cohort can be described as believing *that arthritis is painful*; they share what Loar calls 'social content', which is sensitive to environmental and social factors and is therefore wide. Social content is captured by the content-ascribing *that*-clauses that we use to attribute propositional attitudes. But content-ascribing *that*-clauses do not correspond, one-to-one, to belief contents, as these are individuated in common-sense psychological explanation. So it does not follow that Alf and his cohort have the same belief contents as these are individuated by common-sense psychology.

Loar cites a number of examples in support of his claim. He asks us to consider a diary which we know was written by either an earthling or a Twin Earthling, although we do not know which. One entry says: 'No swimming today; the water is too rough'. We would explain the diarist's behaviour as follows: because she believes that the water is rough, and that if the water is rough one should not swim, she believes that one should not swim. If the explanation is challenged on the grounds that we don't know whether she has any *water* beliefs at all, then, according to Loar, we would likely produce a paraphrase of the original explanation that makes no commitment about the nature of the liquid in question. We might say, for example, 'She believes that the local sample of the liquid that fills the oceans and lakes is rough'. In other words, we would produce a *single* explanation that subsumes the beliefs of both earthlings and Twin Earthlings, rather than two explanations, one adverting

[5] See the discussion in Brown (2003) of Segal's suggestion (2000: 77) that Alf's concepts have undergone a 'natural evolution'.

[6] A radically holist view of concept individuation, according to which *any* change of belief involving the vocable 'arthritis' would count as a change in concept, is, of course, highly controversial. But the view under consideration here is not radically holist. It is not implausible to think that a subject who believes that 'arthritis' refers to a disease of the joints *and* the bones has a different concept.

to water and one to XYZ. On Loar's view, the original explanation attributes to the diarist only a certain way of conceiving of things; in particular, it attributes mental states that are interrelated (and related to stimuli and behaviour) in specified ways, but involves no commitment about the environmental or social context in which the beliefs are embedded, and to which content-ascribing *that*-clauses are sensitive. For the purposes of predicting and explaining behaviour, Loar argues, beliefs are individuated by their role in perception, inference, deliberation, and action. Moreover, this complex role—what is often called 'conceptual role'—determines a kind of content, what Loar calls 'psychological content'.[7] The beliefs that Alf and Twin Alf express by their utterance of the word forms 'arthritis is painful', prior to Alf's correction by his doctor, play the same role in their mental lives; they therefore have the same psychological content. The kind of content relevant for common-sense psychological explanation, Loar concludes, is narrow. We shall return to Loar's claim in the next section when we consider the explanatory role of content attributions.

Noam Chomsky (1992, 1995) has been highly critical of the methodology used to establish the existence of wide content. The externalist thought-experiments, he claims, make illegitimate appeal to speakers' intuitions about such notions as *reference*, *extension*, and *true of*, technical notions about which speakers can have no theory-neutral intuitions (1992: 225; 1995: 42). Moreover, speakers' intuitions about such things as whether a liquid superficially identical to water counts as *the same liquid* are highly sensitive to contextual factors. Change the circumstances or pragmatic presuppositions of the Twin Earth story, he suggests, and judgements will change accordingly. Chomsky doubts that 'any general or useful sense can be given to such technical notions as "wide content" (or any other notion fixing "reference") in any of the externalist interpretations' (1992: 225).

Chomsky is surely right to point out the lability of the intuitions upon which the externalist interpretation of the thought-experiments depends, though his rejection of the whole externalist programme is an extreme and unwarranted reaction. Care must be taken in attributing beliefs to subjects simply on the basis of the conventional meanings of the words they utter.[8] The seemingly innocuous assumption appealed to earlier—that a subject's utterances express her thoughts—does not by itself support the attribution of wide content to the subject's mental states.[9] Moreover, we have seen from the discussion of Putnam's Archimedes example that empirical investigation of a subject's behaviour in a wide range of circumstances is sometimes necessary to determine whether a subject has a natural-kind concept in mind. The investigation will probe just how a subject's judgements are sensitive to contextual factors. But the context-sensitivity of our intuitions and judgements about such matters does not itself undermine the existence of wide content or render the methodology deployed in the thought-experiments problematic.

[7] And what Segal (2000, 2003) calls 'cognitive content'.

[8] This is one moral of Loar (1988), but the point is also made by Bach (1987), Patterson (1990), and Crane (1991).

[9] Chomsky himself is sceptical about the notion of a 'shared public language' and so would reject the idea that the subject's utterances have wide meaning (see Chomsky 1995).

20.3 The Explanatory Role of Content

Various authors have expressed doubts that wide content is well suited to serve the goals of psychological explanation. Barbara Von Eckart questions the legitimacy of at least some types of wide content:

> How *could* certain subtle features of the linguistic practices of persons around you influence the content of your mental representations when those features do not directly impinge on you in any way?
>
> (1993: 264)[10]

The argument presupposes that only features that *causally* affect the subject could be relevant to mental content. By hypothesis, differences in the environmental or social context of twins make no difference to their surface stimulations, and so, the reasoning goes, such features could affect their inner states only by some sort of 'action at a distance' (see, for example, Block 1986: 625; Fodor 1987: 40). At very least, it is claimed, such causal connections as do hold between the environment or social context and intentional agents do not seem to be of the right sort to affect their mental states. But, as Burge and others have noted, the presupposition is false. Contextual or relational properties can affect the *type-individuation* of an object or state without affecting it causally. The birth of a child can make a woman an aunt without the two ever coming into causal contact.

While there is nothing metaphysically suspect about wide content, the underlying worry is that it is nonetheless ill suited to play an explanatory role in psychology. Proponents of narrow content have argued that there is no distinctive explanatory work for wide content to do. Segal (1989) is concerned explicitly with the content attributed in computational theories of cognitive capacities. He argues that David Marr's theory (1982) of early visual processing attributes *narrow* content to the states and structures it posits. The argument depends upon general considerations that, if true, would apply to any account of cognitive mechanisms.[11] Moreover, the intuition underlying the argument is shared by many who think that, whatever the differences in their physical or social environments, twins are *psychologically* identical and should be treated as such by any theory interested in explaining their behaviour or cognitive capacities.

Segal argues that unless subjects differ in their discriminative and recognitional abilities there is no motivation for positing different contents. Describing a Twin Earth scenario where distinct environmental conditions, C1 on earth and C2 on Twin Earth, cause identical internal states, non-intentionally described, to be tokened in twin subjects, he says:

> Ascriptions of representations require top-down motivations. These motivations come from finding some use or purpose to which the content is put. So to motivate the attribution of

[10] Von Eckart raises the worry, but it is not clear that she actually endorses the sentiment expressed in the quote.

[11] I argue in Egan (1995) that Marr's theory attributes *wide* contents to the posited computational structures. Nothing in the present discussion will turn on the interpretation of Marr's theory.

> [distinct environment-specific contents] rather than . . . [narrow contents] we would have to find some purpose that the more specific content could serve, but that the vaguer [narrow] contents could not . . . What would this be? Nothing that issued in any discriminative ability, or recognitional ability that could be detected by any tests of the sort employed by the theorist of vision. For the twins are twins, and will be the same in every testable respect.
>
> (1989: 205)

He goes on to say that positing distinct contents would violate Ockham's razor:

> There would just be no point in invoking the two contents, where one would do. For there would be no theoretical purpose served by distinguishing between the contents.
>
> (1989: 206)

There is, in fact, an important theoretical purpose served by the ascription of wide content in theories of cognitive mechanisms. Spelling this out will require a little stage-setting. I will then consider whether the point carries over to common-sense psychological explanation.

The semantic interpretation of a computational mechanism specifies which properties are tracked by the posited structures when the mechanism is functioning properly in its normal environment. An interpretation of a computational system is given by an interpretation function that specifies a mapping between equivalence classes of physical states of the system—in so-called 'classical' devices these will be symbols or data structures—and elements of some represented domain. To interpret a device as a visual system is to specify a mapping between states of the device and tokenings of visible properties such as changes of depth in the scene. To interpret a device as an adder is to specify a mapping between states of the device and numbers. The specified states of the device are thus construed, under the interpretation, as *representations* of changes in depth or of addends and sums.

Whether a computationally characterized device succeeds in computing, say, the depth of objects and surfaces in the scene from information about the disparity of points in the retinal image depends on whether its internal states co-vary with changes of depth in the environment. This requires a certain *fit* between the mechanism and the world. In the actual world the fit is a product of natural selection. Let us call a visual mechanism adapted to the terrestrial environment 'Visua' and the distal property it tracks—changes of depth—'C1'.[12]

A given computational mechanism would not enhance fitness in every environment. Being an adaptation is a contingent property of a computationally characterized system. We can imagine a duplicate of Visua, Twin Visua, appearing spontaneously (perhaps by random mutation) in a world to which it is not adapted. Imagine that the counterfactual world, E2, is different enough from the actual world that Twin Visua's states track some distal property other than changes in depth. Call this property 'C2'. Since Visua and Twin Visua are physical duplicates, the two mechanisms have the same discriminative and recognitional abilities. Visua would track C2 if it were transported to E2. Twin Visua will contribute to the fitness of the

[12] This is an adaptation of an example from Segal (1989).

organism containing it only if C2 is a useful property to track or represent in E2. C2 is some function of surfaces, local light, and local optical laws, but tracking C2 might not allow the organism containing Twin Visua to recover what is where in the scene. If it does not, we might wonder whether it is even appropriate to call Twin Visua a *visual* mechanism.[13]

The important question for present purposes is 'What do Visua's and Twin Visua's internal states represent?'. It is natural to say that Visua represents C1—changes of depth. It is in virtue of tracking changes in depth in the scene that Visua contributes to the organism's successful interaction with its environment. Perhaps it is also natural to say that Twin Visua represents the distinct property C2. In any case, it would be odd to say that Visua represents some more general property C3 that subsumes both changes of depth (C1) and the strange and (from the terrestrial perspective) unnatural property C2.[14] In other words, there is a clear rationale for attributing to Visua and Twin Visua distinct, environment-specific, wide contents that make apparent the contribution of the mechanism to the success of the organism in its normal environment, rather than an unperspicuous narrow content that does not.

To summarize: Visua and Twin-Visua are *type-identical* computational mechanisms. The computational characterization that specifies a mechanism's basic causal operations subsumes both of them. This environment-general characterization allows us to predict and explain how a mechanism would behave in counterfactual worlds. But it doesn't tell us what the mechanism would *represent* in other environments. The content-determining correlations are those that obtain between states of the device and property tokenings in the local environment. The correlations that obtain in counterfactual environments, where the objects, properties, and laws may be quite different, are not relevant to the semantic interpretation of the mechanism.

I claimed above that it is *natural* to ascribe environment-specific contents to cognitive mechanisms. The reason is not hard to find. The questions that antecedently or pretheoretically define a cognitive theory's domain are typically framed in terms that presuppose the organism's success in its normal environment. We want to know how the organism can recover the depth and orientation of the objects in the scene from information in two-dimensional retinal images.

The cognitive theorist sets out to answer questions that are already framed in environment-specific terms. If the mechanism's states are interpreted as *representing* depth and orientation, rather than more general properties determined by correlations that obtain in counterfactual environments, then the theory is poised to answer these questions.

[13] If Twin Visua is not a visual mechanism, then of course Visua is a visual mechanism only contingently. I argue in Egan (1999) that computationally characterized mechanisms subserve cognitive functions such as vision or parsing only contingently.

[14] An internalist might claim that Twin Visua represents C1, but this seems unmotivated, since Twin Visua's states may never co-vary with C1 in E2. E2 is not a world where Twin Visua (or Visua, for that matter) sees depth.

The semantic interpretation of a computational mechanism can be responsive to these explanatory and pragmatic considerations, addressing the questions that initially motivated the search for and development of the theory, only because a computational characterization provides an environment-independent characterization of the basic operations, or functional architecture, of the device.[15] As noted above, this characterization, which for present purposes we can call a 'formal' characterization,[16] provides the basis for predicting and explaining its behaviour in any environment. The semantic interpretation can serve our more parochial interests. A complete computational account of a cognitive mechanism comprises two independent components, a fully general formal characterization of the causal structure of the mechanism, and an environment-specific semantic interpretation. Computational accounts do not *need* narrow contents to specify the causal structure of the system. The formal component of the theory does that.

Let us return to Visua and Twin Visua. Suppose now that the property tracked by Twin Visua in E2—C2—is a useful property for an organism in E2 to detect. Twin Visua therefore contributes to the fitness of the organism containing it. Imagine that an enthusiastic editor on earth (E1), always on the lookout for new markets, asks the theorist responsible for characterizing Visua to produce a textbook that could be marketed and sold in both E1 and E2. Since the formal component of the theory that specifies Visua's basic causal operations will characterize Twin Visua's as well—Visua and Twin Visua are formally identical mechanisms—the theorist needs only to produce a single semantic interpretation that specifies what this formally characterized mechanism will represent in E1 and E2. Since the mechanism does not track C1 in E2 or C2 in E1, neither *C1* nor *C2* is a plausible candidate. Rather, an interpretation appropriate to both worlds would take the mechanism to represent the more general property C3 that subsumes both C1 and C2.[17]

Notice, first, that the content *C3* is a wide content. The new semantic interpretation specifies what the mechanism represents in E1 and E2, but not what a physically indistinguishable mechanism might represent in some third environment E3. (This follows by an iteration of the reasoning above.) While nonetheless wide, *C3* is, in a sense, *narrower* than either *C1* or *C2*; *C3* prescinds from the environmental differences between E1 and E2. The explanatory interests served by the new interpretation are less local, less *parochial*, than those served by the original interpretation, which was designed to address questions posed in vocabulary appropriate to earth. Whereas the old interpretation answered such pretheoretic questions as 'How is the organism able to recover the depth of objects in the scene' by positing representations of *depth*, the new interpretation provides the basis for answering this question and an analogous question framed in vocabulary appropriate to E2—'How is the organism able to recover information about C2?'—by positing representations of the more general

[15] See Egan (1999) for the argument that computational characterization is *individualist* or narrow.

[16] Following Fodor's usage (1980).

[17] C3 may be understood as the disjunction of C1 and C2, or as a determinable that has C1 and C2 as determinates.

property C3, and supplying auxiliary assumptions about how C3 is related to the locally instantiated properties C1 and C2.

As it happened, the overeager editor was somewhat surprised that sales of the new interplanetary textbook on earth fell off rather sharply from the first edition, designed solely for the local market. Besides introducing a new vocabulary containing such unfamiliar predicates as 'C3', the new edition required cumbersome appendices appropriate to each world, explaining how to recover answers to questions about the organism's performance in its local environment, questions that motivated the search for an explanatory theory in the first place. Readers complained that the new edition was much less 'user-friendly'.

The editor was therefore dissuaded from his original idea of commissioning an intergalactic version of the text, which would provide a genuinely narrow semantic interpretation that would specify what Visua would represent in *any* environment.[18] He came to realize that a semantic interpretation of the basic causal processes of a mechanism is primarily a *gloss* that allows a theory to address local explanatory interests. Any gloss that shows that the theory is doing its job will be couched in a vocabulary that is perspicuous for the local audience with these interests. The moral is that a truly intergalactic cognitive science would not be representational.

So what about common-sense psychological explanation? This sketch of the role of content attribution in cognitive science does not carry over straightforwardly to common-sense psychological explanation primarily because content attribution in our folk practices must do double duty. Nonetheless there are crucial similarities.

Beliefs and desires must be construed as causally efficacious states of subjects to play the predictive and explanatory roles required of them in our folk practices. Whereas computational theories characterize the states they posit in both formal and semantic terms, folk psychology characterizes propositional attitudes only by their contents.[19] Unlike the states involved in the highly modularized, informationally encapsulated processes characterized by our most promising computational theories, propositional attitudes have very complex functional roles which include, typically, accessibility to consciousness. We must rely on content ascriptions to 'get at' the complex functional roles in virtue of which propositional attitudes figure in predictions and explanations of behaviour. Normally content serves this purpose quite well. The state that is typically caused by looking out of a window on a rainy day, and that typically causes (in conjunction with other propositional attitudes) umbrella-carrying behaviour, can be referred to as *the belief that it is raining*. Identifying this state by its content allows us to infer quite a bit about how it

[18] Instead the editor commissioned an environment-specific (wide) semantic interpretation for each world, to accompany the environment-neutral (narrow) formal account of the mechanism.

[19] Proponents of the 'representational/computational theory of mind' (RTM) are hopeful that scientific psychology will eventually provide a formal specification of the complex functional roles of propositional attitudes, and hence will 'vindicate' folk psychology (see Fodor 1987: ch. 1). According to the RTM, propositional attitudes will be revealed as relations to sentences in a language of thought. At present, though, folk psychology must continue to get by without a formal specification of the mechanics of thought.

will interact with environmental conditions and other internal states to produce additional mental states and behaviour.

A moral of Loar (1988) is that content-ascribing *that*-clauses are not perfectly suited to serve the purposes of psychological explanation. Consider the diarist again. According to Loar, we would explain the diarist's behaviour by producing a single explanation that subsumes the beliefs of the diarist and her twin. In this context we don't really care what kind of liquid the subject is thinking about—information that is typically made available by a content-ascribing *that*-clause, such as *that the water is too rough*. What is relevant for predicting and explaining the diarist's behaviour is the causal role of her belief, its pattern of actual and potential interaction with stimuli, behaviour, and other mental states. This narrowly construed causal role, Loar claims, determines a kind of content that captures the way that the subject and her twin conceive of things, considered independently of their environments.

The internalist intuition that the twins are identical in psychologically relevant respects is compelling. But respecting this intuition does not require that they share some narrow content; it requires only that they share some psychological *states*. As Loar himself notes (1988: 105), a 'more minimal' lesson to be drawn from the examples is that the predictive and explanatory schema of common-sense psychology individuate propositional attitudes narrowly, by their (narrowly conceived) causal roles, but not necessarily that causal role determines a second kind of *content*. Since we don't have a way of specifying the shared psychological states of the diarist and her twin directly, in 'intrinsic' terms that spell out precisely the psychologically relevant causal potential of the states, we must find a content ascription that applies to both the diarist and her twin. Such an ascription is not hard to find: the diarist refrained from swimming because she believed *that the ocean (lake, river, or whatever) was too rough*. But there is no reason to think that the content so attributed is narrow. It is, in a sense, *narrower*; the content ascription prescinds from the environment-specific information supplied by the *that*-clause *that the water is too rough*, so it attributes a more minimal environmental commitment. But it does not abstract away entirely from the subject's external surroundings, as it would if it were genuinely narrow. And, just as in the computational case discussed above, it is unlikely that contents that applied intergalactically—to subjects in any environment whatsoever—would serve the complex role required of content. Folk thought and talk about intentional agents are couched in terms that presuppose the local environments in which they are situated. We are not interested simply in the bodily movements of subjects considered in isolation from the objects and properties that make up their world. Wide-content glosses on the (narrowly individuated) internal states that cause intentional behaviour allow environmentally situated creatures such as ourselves to understand these states in terms of their publicly available normal causes and effects.

Content ascribing *that*-clauses—purveyors of wide content—allow us to 'get at' the causal roles, in virtue of which propositional attitudes figure in predictions and explanations of behaviour, only indirectly and imperfectly, as the thought-experiments illustrate. It does not follow that there is a second type of content—so-called 'psychological content' that is supposedly narrow—that does the job perfectly.

Even if we can construct such a notion by various devices, and for various theoretical purposes,[20] there is no reason to think that the notion actually plays a role in our common-sense practices, or that it would serve the other purposes of content as well as ordinary content, which is wide in varying degrees. We can use a wide content ascription to 'sneak up on' the narrowly individuated causal role shared by the diarist and her twin, or Oscar and Twin Oscar, or Alf and Twin Alf, withdrawing to a more minimally committed, but still wide, content when the occasion requires.

The diarist example indicates that wide content can be too fine-grained to capture what is relevant for psychological explanation. A second example from Loar (1988) suggests that it can be too coarse-grained. Imagine that, some months before we encounter him in Burge's story, Alf visits Paris and hears of a disease called 'arthrite' that affects the joints. Not realizing that 'arthritis' and 'arthrite' are intertranslatable, and being a bit of a hypochondriac, Alf comes to believe that he has two problems with his ankles, one called 'arthritis', which he later comes to believe has spread to his thigh, and a second ailment called 'arthrite'. Loar claims that Alf has *two* beliefs with distinct causal powers. Alf's 'French' belief that he has arthritis in his ankles does not interact in the appropriate way with his 'English' belief that if he has arthritis in his ankles he should take aspirin, although his 'English' belief that he has arthritis in his ankles does. There are two distinct states with different interactive potential, but only one belief ascription—*the belief that he has arthritis in his ankles*—available to pick out the distinct states. Loar concludes that Alf's 'arthritis' and 'arthrite' beliefs have distinct 'psychological' contents, determined by their distinct causal roles.

Once again, a more minimal conclusion is preferable. Alf does seem to have two distinct states with different causal roles. It is not hard to find distinct content-ascribing *that*-clauses that characterize the difference in the causal roles and hence serve to pick out the underlying states. We might say 'He thinks that the original problem with his ankles can also occur in muscles and bones, but that the new disorder affects only the joints'. There is no reason to think that the content ascribed thereby is *narrow*. We can say, with Loar, that the distinct content ascriptions capture 'the way the subject himself conceives of things' (1988: 108). The subject conceives of himself as having two distinct problems with his ankles. But 'the way the subject conceives of things' is not a name for narrow content.

To summarize. Content plays a role in our common-sense explanatory practices that it does not play in computational accounts of cognitive mechanisms. It allows us to refer to the causal roles in virtue of which propositional attitudes figure in predictions and explanations of behaviour.[21] This is the primary job of content ascription in common-sense, as opposed to scientific, psychology, but it is a job that *that*-clauses—purveyors of wide content—do only imperfectly. When necessary, as in the examples discussed by Loar, we can refine our practice, relying on more elaborate *that*-clauses that reveal aspects of the causal roles of the underlying states that are

[20] Chalmers's *epistemic content* (2002) is one proposal.

[21] The formal component of computational theories, rather than the content component, specifies the relevant functional architecture of the system.

relevant for prediction and explanation. But the content attributed when we try to characterize the way the subject conceives of things is still wide.

Content ascription serves other purposes in our common-sense practices besides enabling the prediction and explanation of behaviour. Some of these roles exploit the environmental or contextual specificity that is the hallmark of wide content. I shall mention two.

As several authors have noted (see, for example, Field 1978; Loar 1981, 1988; Schiffer 1981), we make use of the fact that people's beliefs are often reliable indicators of the way the world is. As Field (1978) puts it:

> [W]e are constantly using our opinions about other people's beliefs in forming opinions about the world. The fact that a child believes that he has done something I won't like (a fact that can often be inferred from his behaviour) gives good reason to think he *has* done something I won't like; the fact that most physicists believe that there are gravitational waves . . . is good reason for me to believe that there *are* gravitational waves; and so forth.
>
> (1978: 103)

On occasion, when the point of a belief ascription is to convey information, a speaker might ascribe a content that she knows does not capture the way the subject conceives of things. One might say, for example, that Perry White, the editor of the Daily Planet, thinks that Superman has real literary talent. White, of course, does not know that his star reporter, Clark Kent, is Superman, but the belief ascription serves its purpose—conveying information about Clark Kent/Superman—as long as the audience knows that Clark Kent is Superman.

Burge (1978) identifies another use of content ascribing *that*-clauses. They constitute a *standard* by reference to which a subject's grasp of the commitments of public discourse may be evaluated, and her intellectual responsibilities with respect to that discourse assessed. Thus, we can hold an individual answerable to the community standard, and subjects can hold themselves so answerable, even when they have only incompletely mastered the elements within the standard. Alf, in Burge's original thought-experiment, can take himself to have been *wrong* about arthritis, and to have been corrected by his doctor.

This purpose of content attribution—providing a normative standard against which individual behaviour can be evaluated—can be at cross purposes with the goal of predicting and explaining behaviour. It is this tension that gives rise to the dispute over the correct interpretation of Burge's example. Prediction and explanation will fail when the subject falls short of the standard, if we read the subject's mental states directly off his words. If we want to predict and explain the subject's behaviour, we must look for a distinct, and typically more elaborate, *that*-clause that better captures the causal role of the subject's psychological state. As Loar would put it, we try to characterize 'the way the subject himself conceives of things'. But, I have argued, the content attributed is still wide.

Notice that these two uses of content presuppose *local* interests. The information conveyed by a belief attribution is often about the local environment. In the above case, it is about *Superman*, not his twin in some other world. A subject's competence

will be evaluated by the standards of the local linguistic community. When the goal of content ascription is the prediction and explanation of behaviour, I have argued, content ascribing *that*-clauses are used to pick out the causal roles of our beliefs and desires, properties shared by twins. All of these purposes are served by wide content. We are sometimes interested explicitly in how the subject would behave if the world were different in various ways, and the relevant *that*-clauses may prescind from some features of the environment, as they do in the diarist case. So content can be wide in varying degrees. Genuinely narrow content, content that supervenes on the intrinsic states of the subject and so would apply in any environment, is an idealization along this dimension. But it is too general and too abstract to bear an explanatory burden.[22]

References

Bach, K. (1987), *Thought and Reference* (Oxford: Oxford University Press).

Block, N. (1986), 'Advertisement for a Semantics for Psychology', *Midwest Studies in Philosophy*, 10: 615–78.

Brown, J. (1998), 'Natural Kind Terms and Recognitional Capacities', *Mind*, 107: 275–303.

____ (2003), 'Externalism and the Fregean Tradition', in A. Barber (ed.), *Epistemology of Language* (Oxford: Oxford University Press), 431–55.

Burge, T. (1978), 'Individualism and the Mental', *Midwest Studies in Philosophy*, 4: 73–121.

Chalmers, D. (2002), 'The Components of Content', in Chalmers, *Philosophy of Mind: Classical and Contemporary Readings* (Oxford: Oxford University Press), 608–33.

Chomsky, N. (1992), 'Explaining Language Use', *Philosophical Topics*, 20: 205–31.

____ (1995), 'Language and Nature', *Mind*, 104: 1–61.

Crane, T. (1991), 'All the Difference in the World', *Philosophical Quarterly*, 41: 1–26.

Egan, F. (1995), 'Computation and Content', *Philosophical Review*, 104: 181–203.

____ (1999), 'In Defense of Narrow Mindedness', *Mind and Language*, 14: 177–94.

Evans, G. (1982), *The Varieties of Reference* (Oxford: Clarendon).

Field, H. (1978), 'Mental Representation', *Erkenntnis*, 13: 9–61; repr. in N. Block (ed.), *Readings in the Philosophy of Psychology*, ii (Cambridge, Mass.: Harvard University Press, 1981), 78–114. Page references are to the reprint.

Fodor, J. (1980), 'Methodological Solipsism Considered as a Research Strategy in Cognitive science', *Behavioral and Brain Sciences*, 3: 63–109.

____ (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Cambridge, Mass.: MIT Press).

Loar, B. (1981), *Mind and Meaning* (Cambridge: Cambridge University Press).

____ (1988), 'Social Content and Psychological Content', in R. Grimm and D. Merrill (eds.), *Contents of Thought* (Tucson, Ariz.: University of Arizona Press), 99–110.

Marr, D. (1982), *Vision* (New York: Freeman).

Patterson, S. (1990), 'The Explanatory Role of Belief Ascriptions', *Philosophical Studies*, 59: 313–32.

Putnam, H. (1975), 'The Meaning of "Meaning"', in Putnam, *Mind, Language, and Reality: Philosophical Papers*, ii (Cambridge: Cambridge University Press), 215–71.

[22] Thanks to Brian McLaughlin and Robert Matthews for helpful comments on an earlier version of this paper.

Schiffer, S. (1981), 'Truth and the Theory of Content', in H. Parret and J. Bouverese (eds.), *Meaning and Understanding* (Berlin/New York: de Gruyter).
Segal, G. (1989), 'On Seeing What Is Not There', *Philosophical Review*, 98: 189–214.
——(2000), *A Slim Book about Narrow Content* (Cambridge, Mass.: MIT Press).
——(2003), 'Ignorance of Meaning', in A. Barber (ed.), *Epistemology of Language* (Oxford: Oxford University Press), 415–30.
Von Eckart, B. (1993), *What Is Cognitive Science?* (Cambridge, Mass.: MIT Press).