# *Individualism and Perceptual Content*

MARTIN DAVIES

Tyler Burge (1986) has argued that the practice of psychology—particularly the practice of the psychology of vision—is not individualistic. This is to say, roughly, that the semantic contents of states cited in psychological theory do not supervene upon the internal constitution—physical, chemical, neural, or functional—of the creature whose states they are. Not everyone has been convinced by Burge's arguments, and some have expressed their scepticism in print (e.g. Matthews 1988; Segal 1989; McGinn 1989). Nevertheless, I claim, Burge is right; in the following pages, I explain why.

In fact, Burge's paper has two main phases. In the first phase, Burge rejects arguments (e g. Fodor 1986 and 1987, Chapter 2) that purport to show that any psychology worthy of the name "science" *must* restrict itself to individualistic taxonomies. Then, in the second phase, Burge argues that psychology, as we actually find it, is not individualistic, and that there are powerful, general, reasons to reject individualism about perceptual content.

I shall not focus at all on the first phase. I agree with Burge that scientific psychology is not obliged to restrict itself to individualistic taxonomies (Davies 1986; Jackson and Pettit 1988; McGinn 1989, Chapter 2); in particular, it is not obliged to restrict itself to an individualistic notion of intentional content. In fact, I am even sceptical as to whether there is such a thing as individualistic, or narrow, content; but that is not the issue here. Even supposing that a notion of narrow content is available, if psychology also employs a non-individualistic notion of content—or, say, a spectrum of broader and narrower notions—then psychology is not individualistic.

My concern is exclusively with the second phase of Burge's paper. In §2 and §3, I review Burge's two main lines of argument. After those two sections, the conclusion is that the most promising line of attack for the individualist is to be somewhat revisionary about the attribution of perceptual content. In §4 and §5, I go on to examine, and ultimately to reject, a revisionary individualist strategy that is suggested by both Robert Matthews (1988) and Gabriel Segal (1989). In §6, I identify what I take to be the error in the revisionary strategy. Thus, I claim, Burge is correct: the content of perceptual experience is not individualistic.

## *1. Preliminary clarification*

Before starting upon the main business, however, I offer some brief clarification of the key notions of perceptual content and individualism.

**Mind, Vol. C . 4 . October 1991**
**0026-4423/91 $3.00**

*1.1 Perceptual content*

Perceptual content is a kind of *non-conceptual content*. If a subject has a belief with a certain content, then the subject must grasp the constituent concepts of that content: belief content is conceptualised content. In contrast, a subject may have an experience without possessing the concepts that would be used in the specification of the content of that experience (Evans 1982, pp. 151–70). For example, in order to undergo an experience that presents the world as containing a cube four feet in front of her, a subject does not need to employ the concept of a cube; indeed, she does not even need to possess the concept of a cube.

The subject will, of course, employ the concept of a cube if she takes the experience at face value and judges—and so believes—that there is a cube four feet in front of her. The content of her judgement is conceptualised content; and we can say that the experience that prompts the judgement has that conceptualised content too. The point is not to deny that experiences have conceptualised content, but to affirm that experiences have a kind of content that is representational though non-conceptual.

It is a controversial question whether conceptual content is essentially tied to the use of public language. But, since perceptual content is a kind of non-conceptual content, it seems reasonable to suppose that perceptual content is "fully independent of [public] language" (Burge 1986, p. 26; Burge is explicitly considering "the representations of early vision"). The experiences of creatures that lack language may still have perceptual content; and the completely general claim that perceptual content is individualistic can be probed by considering such creatures.

Finally, it is natural to suppose that perceptual content is *not object-involving*. It is plausible that, if two objects are genuinely indistinguishable for a subject, then a visual experience of the one has the same perceptual content as a visual experience of the other. But even if this view about perceptual content is not ultimately correct, we can focus our attention upon the existentially quantified content of experiences. For example, a visual experience may present the world as containing *a* cube four feet in front of the subject. It matters not at all to that (existentially quantified) perceptual content of the subject's experience which cube it is that she is looking at.

Because perceptual content is non-conceptual content, is independent of language, and is not object-involving, familiar anti-individualist arguments about belief content (e.g. Putnam 1975; Burge 1979; McGinn 1982) do not automatically transpose into anti-individualist arguments about perceptual content. Thus, the possibility is opened up that perceptual content is both representational and individualistic. Perceptual content is representational, in that the content of an experience is specified in terms of how the external world would have to be for the experience to be veridical. But perhaps perceptual content is also individualistic, in that any duplicate of the perceiving subject would be having an experience with just the same content. It is this possibility that Burge aims to close off.

*1.2 Individualism*

I have been equating individualism about perceptual content with the claim that perceptual content supervenes upon internal constitution, so that it is preserved across duplicates. Strictly speaking, however, the claims are distinct. Individualism is a claim about individuation. In the present context, individualism is the claim that

> there is no necessary or deep individuative relation between the individual's being in states of those [mental] kinds and the nature of the individual's physical or social environments. (Burge 1986, p. 4)

So construed, individualism entails the claim of *local supervenience*: the mental kind of an individual's mental state could not be different unless the individual's internal constitution were different.

But the entailment does not quite hold in the opposite direction: anti-individualism is consistent with local supervenience. For suppose that, as a matter of metaphysical necessity, an individual's environment could not be different unless the individual's internal constitution were also different. Then sameness of internal constitution would suffice for sameness of mental kind, even if the mental kind depended for its individuation upon the external environment.

The upshot is that a failure of local supervenience is sufficient to establish anti-individualism; but anti-individualism does not strictly guarantee any failure of local supervenience. What I shall be arguing here is that perceptual content is not locally supervenient; and I shall use the term "anti-individualism" loosely, to include the claim that local supervenience fails.

Finally, we need to distinguish between a *conservative* and a *revisionary* individualist stance towards an example of perceptual content. Consider, for example, such a workaday perceptual content as that there is a cube four feet in front of the subject. For the purposes of this paper, to adopt a conservative individualist stance towards this example is to accept that specification of content, and to argue that the content is locally supervenient. In contrast, to adopt a revisionary individualist stance towards the example is to say that, in full seriousness and for scientific purposes, that workaday content should be replaced by a content that is locally supervenient.

## 2. *The first argument: Psychological practice*

Burge offers two main arguments against individualism in psychology. One (1986, §II, and especially p. 34) aims to show that the actual practice of psychology is not individualistic. The other (1986, §III, and especially p. 41; cf. Burge 1988a) abstracts from any details of psychological theory, and sets out to demonstrate that perceptual content does not even supervene upon internal constitution *plus* behavioural dispositions. In this Section, I focus upon the first argument.

Burge's first argument (1986, p. 34) is based upon three points that are illustrated from Marr's theory of vision. These are:

> [T]he theory makes essential reference to the subject's distal stimuli and makes essential assumptions about contingent facts regarding the subject's physical environment. (p. 29)
>
> [T]he theory is set up to explain the reliability of a great variety of processes and sub-processes for acquiring information, at least to the extent that they are reliable. (p. 29)
>
> [T]he information carried by representations—their intentional content—is individuated in terms of the specific distal causal antecedents in the physical world that the information is about and that the representations normally apply to. (p. 32)

The conclusion of the argument is that "individualism is not true for the theory of vision" (p. 34).

Now, this conclusion about individualism strictly so-called is immediate given the premise that the content of visual representations "is *individuated* in terms of the specific distal causal antecedents in the physical world ..." (p. 32; my emphasis). But Burge's actual argument proceeds to that conclusion via the failure of local supervenience. The crucial staging post along the way is the following claim:

> if these physical conditions [by which the representations are normally caused and to which they normally apply] and, possibly, attendant physical laws were regularly different, the information conveyed to the subject and the intentional content of his or her visual representations would be different. (p. 34)

This is step three in a six-step argument, and it has to be acknowledged that the defence of the six steps does not leave much room for objections to be entered.

The tightness of the argument is unsurprising. For Burge invokes the further assumption that

> There is no metaphysically necessary relation between individualistically individuated processes in a person's body and the causal antecedents of those processes in the surrounding world. (p. 35)

And, as we have already seen (§1.2), if we have that assumption then there is no distance at all from the falsity of individualism strictly so-called (guaranteed by the premise about individuation) to the consequence that perceptual content is not locally supervenient.

*2.1 The individualist's response. Guides and theories*

Burge's arguments address the conservative, rather than the revisionary, individualist. That is, Burge takes it for granted that a subject's visual experiences have such contents as that there is a cube four feet in front of her. Since there is scarcely daylight between the steps of Burge's first argument, the only feasible line of response for a conservative individualist is to mount an attack upon the premise about individuation. The conservative individualist's strategy must be to deny that it is implicit in Marr's theory of vision that there is a "necessary or deep individuative relation" (1986, p.4) between an experience's having a content concerning cubes, for example, and facts about the "specific distal causal antecedents in the physical world" (p. 32).

The individualist can prepare the ground for such an attack by noting that it is typical of the scientific study of cognitive processes that it makes use of the notion of representation without itself offering any constitutive theory of what it is for a state to have a particular representational content. Thus, for example, we find Paul Smolensky saying:

> I will not address the question: aside from the modeler's say-so, what grounds do we have for believing that the units in question really do represent what they are claimed to? This is an important open philosophical question about the nature of representation in the connectionist approach. (1990, p. 131)

As in the case of computational models, so also in the case of the empirical study of actual human cognitive processes, science does not wait for a satisfactory philosophical answer to the important open question about the nature of representation.

However, even without a constitutive theory to hand, a psychologist can quite reasonably suppose that certain states have particular contents. Consider, for example, the kind of visual experience produced in normal conditions by a cube four feet in front of the subject—a kind of experience that reliably covaries with the presence of a cube, and is the result of the operation of a mechanism that has evolved to deliver information about the shapes of distal objects, a kind of experience that produces behavioural consequences appropriate to the presence of a cube, a kind of experience which, when taken at face value, leads to the judgement that there is a cube, and perhaps to the utterance of a public language word that means *cube*. It would certainly be surprising if a theory were to pronounce that this kind of experience does not have the content—perhaps *inter alia*—that there is a cube four feet in front of the subject.

The cited properties of the experience type surely provide a good *heuristic guide* to the presence of a certain content; but this tells us relatively little about the optimal shape for a *philosophical theory* of content. A theory of perceptual content might attach weight to causal antecedents, to teleological purpose, or to behavioural consequences; or it might take as given the notion of the content of judgement—or, with less plausibility, the content of public language utterances—and analyse the notion of perceptual content in terms of the way in which experience leads to judgement—or even to utterance.

So, if we find a psychologist offering grounds for the attribution of certain contents to states, it is legitimate to ask whether he is sketching a philosophical theory, or merely employing heuristic guides. Segal registers a sharp disagreement with Burge on just this point:

> Burge claims that Marr uses such items to individuate contents. In my view, he merely uses them as a guide to discovering what the contents are. (1989, p. 199)

In fact, Segal complains:

> The conflation of what is taken to justify the ascription of a given content with what makes it the case that the content is there, is apparent at numerous points in Burge's exposition. (1989, p. 212)

*2.2 The individualist's response: Counterfactual cases*

With the ground thus prepared, the conservative individualist may now attack Burge's first argument, by arguing that psychological practice does not force an anti-individualist verdict upon the crucial counterfactual cases.

In order to test the individualist thesis, we imagine a creature to which Marr's theory actually applies—or a duplicate of such a creature—existing in a different physical environment. In this imagined environment, an intrinsic state that is representational is apt to have "specific distal causal antecedents in the physical world" that are different from those in our actual situation. We now ask: Are the contents of this creature's representations in the imagined situation the same as, or different from, the contents of the creature's representations in the actual situation?

If the answer is that the contents are different across the two circumstances, then this demonstrates the anti-individualist thesis that perceptual content is not generally preserved across duplicates:

> A person's intentional states and events could (counterfactually) vary, even as the individual's physical, functional (and perhaps phenomenological) history, specified non-intentionally and individualistically, is held constant. (Burge 1986, p. 4)

If, on the other hand, the answer is that the contents are preserved across these two circumstances which differ only in aspects external to the creature, then the thought experiment does nothing to unseat individualism.

The individualist argues that the first answer is not obligatory. We can certainly agree that psychologists of vision speak of stages of processing leading from a grey level array to a representation, say, of a sphere of such-and-such a size at such-and-such a distance. Even if the physical laws are held constant, the facts about the grey level array seriously underdetermine the nature of the distal stimulus; so if we think of the processing as inferential then the inference is enthymemic. The missing premises take the form of large scale assumptions about the physical environment. Because the information processing system builds in those assumptions, it yields determinate representations of distal objects. Because the assumptions are largely true, the resulting representations are largely veridical. And it is no accident that the assumptions are true, given the adaptive advantage to the creature of veridically representing its environment.

All this can be accepted, the individualist continues, but from it nothing much follows about our creature in the imagined counterfactual situation. It is open to us to suppose that the creature is doomed to be the victim of a good deal of misrepresentation. This verdict on an imagined counterfactual case is in no way inconsistent with "the success-orientation of [Marr's] theory" (p. 34), because success-orientation governs the methodology of the theory as it is applied to creatures in the actual world. If we now imagine one of those creatures, or a duplicate, in a quite different and inhospitable environment—an environment to which the creature's visual processing system is not adapted—then the methodology does

not require that, even in those counterfactual circumstances, success should be the general rule.

Confident that he has rebutted Burge's first argument, the individualist might offer a variety of diagnoses of what he regards as Burge's mistake. It might be suggested, for example, that Burge's reading of psychological practice follows too closely the contours of a causal theory of reference for expressions of a public language. Thus, Segal charges:

> Burge's argument that the computational theory is not individualistic depends upon reading into it a particular form of the causal theory of reference. A causal theory of that form might be true of some parts of natural language and even of some parts of folk psychology, but one cannot automatically assume that it applies to a computational theory of a perceptual module. (Segal 1989, p. 203)

But these putative diagnoses of Burge's supposed error are of limited interest, since the individualist's attack misses its mark. It is true that Burge does not demonstrate that the contents of perceptual experiences would inevitably be different if the distal causes of those experiences were different. But that is not what needs to be demonstrated in order to establish anti-individualism. We can put the same point another way by saying that individualism does not follow from the fact that there are *some* environmental changes under which perceptual content would be preserved across duplicates.

*2.3 The force of the first argument*

Let us return to that third point about Marr's theory of vision—the premise about individualism:

> [T]he information carried by representations—their intentional content—is individuated in terms of the specific distal causal antecedents in the physical world that the information is about and that the representations normally apply to. (Burge 1986, p. 32)

The word "normally" here is not supposed to be idle. On the contrary, a great deal is packed into it. This is clear from Burge's own description of the kind of counterfactual situation against which the individualist thesis should be tested:

> If [the] theory were confronted with a species of organism reliably and successfully interacting with a different set of objective visible properties, the representational types that the theory would attribute to the organism would be different, regardless of whether an individual organism's physical mechanisms were different. (p. 34)

And again:

> If the environmental conditions were different, the same proximal *visual* stimulations could have regularly had different distal causes. In principle, we can conceive of some regular variation in the distal causes of perceptual impressions with no variation in the person's individualistically specified physical processes, even while conceiving the person as *well adapted* to the relevant environment .... (pp. 35–6)

Consequently, taking a creature to which Marr's theory actually applies and imagining that creature—or a duplicate—in a different and inhospitable environment is not the relevant test for individualism.

Rather, what Burge imagines is that the creature in the counterfactual situation is as well-adapted to its environment as we are to our actual environment. Psychologists following Marr's success-oriented methodology are imagined to set about studying the members of this creature's species. And the claim is that the content-using descriptions employed by those imagined psychologists may well differ from the content-using descriptions employed by actual psychologists, even though the creature being described is, by happenstance, a duplicate of an actual creature. The individualist's attack does nothing to unseat this anti-individualist claim.

The conservative individualist has limited room for manoeuvre, at this point. One possible move is to accept that perceptual content is externally individuated, but to deny that, strictly speaking, psychology types perceptual states by their content. In its functionalist rendition, for example, this strategy would say, first, that having an experience with a certain perceptual content is to be identified with being in a certain role state, and second, that strictly speaking psychology classifies experiences in terms of realizer states rather than role states (cf. Jackson and Pettit 1988). I shall not discuss this strategy further here. Suffice it to say that it is not a strategy that is clearly motivated by actual psychological practice. It seems, rather, to be a strategy that would only be motivated by an antecedent argument for the view that psychology *must* restrict itself to individualistic classification.

The most promising move for the individualist is to depart from conservatism by adopting a revisionary stance towards some examples of perceptual content. Different individualists may opt for different mixtures of conservatism and revision. Elsewhere (Davies 1992), I argue specifically that perceptual contents involving shape and distance properties are not locally supervenient. Consequently, I would claim that an individualist is obliged to adopt a revisionary stance towards such contents—to maintain, for example, that for scientific purposes no experience should really be described as presenting the world to a subject as containing a cube four feet in front of her. I shall consider a revisionary individualist strategy (suggested by Mathews 1988 and Segal 1989) after introducing Burge's second argument.

## *3. Second argument: Objectivity and error*

Burge's second argument is intended to show "that a person's non-intentional dispositions could not fix (individuate) the intentional content of the person's visual presentations" (1986, p. 39). (The argument also appears in Burge 1988a, pp. 69–76.) The key idea from which this argument begins is that visual experiences represent objective states of affairs, so that there is an "is"/"seems" distinction.

Thus, any account of the content of such experiences must make room for the possibility of error, that is, of misrepresentation.

The argument involves two visible, objective, properties O and C. In the particular instantiation of the argument that Burge gives later (1986, pp. 42–3; 1988a, pp. 75–6) the two properties are being a shadow of a certain size and shape (O) and being a similarly sized crack (C). The argument is intended to be general; but the instantiation certainly improves intelligibility; so I shall enter the instantiation parenthetically throughout the discussion.

We suppose that as things actually are a person P often sees occurrences of O (shadows) as O (as shadows). Since we acknowledge the possibility of misrepresentation, we allow that a representation of the type that is usually caused by occurrences of O (shadows) may on occasion be produced in P by an occurrence of the different visual property C (a crack). On such an occasion, P sees an instance of C (a crack) as an instance of O (as a shadow). Furthermore, we may suppose that, on such an occasion, the subject's dispositions to behaviour consequent upon the occurrence of the experience do not furnish the subject with the ability to discriminate the instance of C (the crack) from instances of O (from shadows).

Now, we construct a counterfactual scenario. First, we hold fixed P's internal constitution and behavioural dispositions. Second, we alter the environmental set-up so that there are no occurrences of property O (no shadows of that size and shape). Third, we also alter the external environment so that experiences of the intrinsic type normally produced by occurrences of property O (shadows) in the actual set-up are now normally produced by occurrences of property C (cracks).

Finally, we imagine that, on a particular occasion, P has a visual experience caused by an occurrence of C (a crack), in a way that is normal for this counterfactual scenario. Consequently, the experience produced on this occasion is intrinsically just the same as the experience produced in the actual situation when P misperceives a C (a crack) as an O (as a shadow)

Burge makes two claims about this imagined situation (1986, p. 41). First: "P's visual representation ... would not be of intentional type O′". Here, the intentional type O′ is the type of the experience normally produced by O in the actual situation. So the claim is that in the counterfactual situation, P does not see the occurrence of C (the crack) *as* an O (as a shadow). Second: "P may counterfactually be perceiving something (say, a C) correctly (as a C)—if the processes that lead to that visual impression are normal and of the type that normally produces the visual impression that P has on that occasion". In sum, then, "he or she perceives a C as a C, not as an O" (1988a, p. 74).

If these two claims are correct, then Burge's second argument surely establishes his anti-individualist position.

*3.1 The individualist's response: Covariance theories of content*

Both claims are defended in terms of a premise that is summarised as follows:

> [S]ome visual representations that represent objective entities as such must have the representational characteristics that they have partly

> *because* instances regularly enter into certain relations with those objective entities. (Burge 1986, p. 40)

On a very cursory and superficial reading (ignoring, for example, the word "partly"), it may appear that Burge here relies upon a causal covariance theory of perceptual content. Given such a theory, the first claim follows: P's visual experience cannot present the occurrence of C (the crack) as O (as a shadow), because in the counterfactual situation there are no instances of property O (no shadows of that size and shape) for that type of experience to covary with. Similarly, given such a theory, the second claim follows: P's visual experience can present the occurrence of C (the crack) as C (as a crack), because in the counterfactual situation that type of experience normally covaries with occurrences of C (with cracks).

This interpretation of Burge's second argument immediately invites several responses on behalf of the individualist. First, to the extent that Burge is committed to a covariance theory of content there is a tension within his own position. For the argument begins from the possibility of misrepresentation; yet causal covariance theories notoriously have problems allowing for misrepresentation (Dretske 1986).

Second, to the extent that the argument relies upon a covariance theory of content, it is cast into doubt by a certain type of example. We imagine the actual situation just as in Burge's scheme, and we imagine that in the counterfactual situation the experience type that actually covaries with O (with shadows) covaries instead with C (with cracks), but still presents those occurrences of C (those cracks) *as* O (*as* shadows)

Matthews presents an example of just this kind:

> Suppose... that the shadows and cracks in question are important to the organism's adaptive success, e.g., that the shadows are important sources of shade for the organism during the heat of the day, and that the cracks are large enough that the organism risks injury if it should fall into them. If in the counterfactual environment the organism repeatedly fell into the cracks when during the heat of the day it sought shelter from the sun, we would surely conclude that in this environment the organism perceives cracks as shadows, or at least not as cracks. We would do this, even though the organism's perceptual representations were in that environment "normally" caused or occasioned by cracks. (1988, p. 83)

And he draws the conclusion that there is "something amiss" with Burge's argument.

But the reading of Burge that invites these responses does not do full justice to his position. Burge is not committed to saying that causal covariance is a necessary condition for perceptual content. For at the very point where he might seem to rely on a simple covariance theory of content, he actually introduces the idea of evolutionary history. In a footnote to the passage quoted above, he says:

> Some of the interaction that leads to the formation and representational characters of certain innate perceptual tendencies (or perhaps even representations) may occur in the making of the species, not in the learning histories of individuals. Clearly this complication could be incorporated

> into a generalization of the ... premise—without affecting the anti-individualistic thrust of the argument. (1986, p. 40, fn. 22)

*A fortiori*, Burge is not committed to saying that causal covariance is a sufficient condition for perceptual content. An unfortunate creature might spend its whole life misperceiving cracks as shadows; just as an unfortunate laboratory frog might spend its whole life mistaking beebees for flies.

*3.2 The force of the second argument*

If a causal covariance theory were correct for perceptual content, then certainly perceptual content would not be locally supervenient. In fact, because causal covariance theories of content are exclusively input-side theories, they nowhere advert to output factors such as behavioural dispositions. So, if we take it as a premise that a causal covariance theory of content is correct, then we can show that content does not even supervene upon internal constitution plus behavioural dispositions. There is no doubt that, given that premise, Burge's second argument goes through.

We have just seen that Burge is not committed to a covariance theory of perceptual content. But that does not threaten the second argument since, of course, the argument does not actually require the acceptance of such a theory as a premise. In fact, it is worth reflecting for a moment on just how little the dialectical situation requires of Burge.

The individualist claims that, whatever the differences in environment, identity of internal constitution is a sufficient condition for sameness of content. All that needs to be the case for the individualist to be wrong is that there are some environmental differences—however thoroughgoing—that suffice for a difference of content across duplicates. It is not required that a mere difference in the distal causes of visual experiences should induce a difference of content. Rather, it is enough that there should be some way of fleshing out Burge's admittedly rather underdescribed examples in order to occasion such a difference of content. (And to show what the second argument explicitly sets out to show, there would have to be a way of fleshing out the examples which also preserved behavioural dispositions across the actual and counterfactual situations.)

Nothing in the individualist's response suggests that this requirement cannot be met (see Davies 1992). Once again, the conservative individualist has very little room for manœuvre.

*3.3 A revisionary suggestion*

As I noted at the end of §2, the most promising move for the individualist is to take a few steps away from conservatism. The individualist can adopt a revisionary stance towards the example of the shadows and cracks, and reject the conservative attributions of content to the creature in the actual situation.

Burge makes a good deal of the fact that his argument does not actually require that P perceives occurrences of C (cracks) as C (as cracks) in the counterfactual situation. All that is needed for a difference of content between the actual and

counterfactual situations is that, in the counterfactual situation, P should not perceive occurrences of C (cracks) as O (as shadows). Thus, although he says:

> *On such counterfactual occasions*, P *would be visually representing* small cracks as small cracks. ... Counterfactually, P correctly sees the cracks as cracks (1986, p. 43)

in the later discussion of the argument Burge is explicit that

> nothing in the argument depends on attributing any specific perceptual states to the organism in the counterfactual situation. All that is important is that it be plausible that the counterfactual perceptual states are different from those in the actual situation. (1988b, p. 95)

So far as the requirements of the argument go, it would be enough if, in the counterfactual situation, P perceives occurrences of C (cracks) as, for example, O *or* C (as shadows or cracks; or better, as shadows-or-cracks). This would be enough, given—what the conservative individualist allows—that, in the actual situation, P perceives occurrences of O (shadows) as O (as shadows) and on occasion misperceives an occurrence of C (a crack) as O (as a shadow).

But now the individualist who is prepared to be revisionary may urge that, even in the actual situation, P perceives occurrences of O (shadows) as shadows-or-cracks. This is what Matthews suggests:

> An organism may perceive O's in the actual environment and C's in the counterfactual environment, not as O's or [as] C's, but rather as instances of an objective type that includes both O's and C's. (1988, p. 83)

It is time to consider this revisionary line of individualist argument in more detail.

## *4. A revisionary strategy: Disjunctive contents*

The line of argument that is briefly suggested by Matthews is developed in much more detail by Segal (1989), who uses it to combat Burge's first argument just as much as the second. Recall Segal's remark that we quoted earlier:

> Burge's argument that the computational theory is not individualistic depends upon reading into it a particular form of the causal theory of reference. A causal theory of that form might be true of some parts of natural language and even of some parts of folk psychology, but one cannot automatically assume that it applies to a computational theory of a perceptual module. (p. 203)

Segal goes on to develop this hint that folk psychology and scientific psychology may differ in their attributions of content. The idea is, of course, familiar from discussions of narrow content (Fodor 1986 and 1987, Chapter 2). There it surfaces as the claim that, although folk psychology attributes broad contents to mental states such as beliefs, a properly scientific psychology will instead employ a locally supervenient notion of narrow content. Applied to the case of perceptual content, the idea is that a folk psychological description of a visual experience might say that it presents the world to a subject as containing a cube four feet in

front of her, or as containing a shadow of a certain size and shape, for example. But, to the extent that the description is not individualistic—to the extent that Burge can construct anti-individualistic examples for a particular kind of content attribution—that just shows that science would not really use such a description of the experience. Instead, a scientific psychology would describe the experience in terms of a disjunctive content.

This is a revisionary move, given the way that we have characterised the distinction between conservative and revisionary individualist stances. But it retains one important feature of a conservative stance in that the content recommended for scientific employment is fully representational: it determines how the world has to be for the experience to be veridical.

What Matthews and Segal offer is a general revisionary strategy for responding to putative anti-individualist examples, by attributing disjunctive contents to perceptual experiences. Thus, where Burge claims to have provided an example in which a creature sees an occurrence of C (a crack) *as O (as a shadow)* and a duplicate sees an occurrence of C (a crack) *as C (as a crack)*, the advocate of the strategy argues that both see an occurrence of C (a crack) *as O-or-C (as a shadow-or-crack)*.

The use of the term "disjunctive contents", and the presentation of the strategy in terms of hyphenated disjunctions is mine and not Segal's own. What Segal actually says of the strategy is that

> we ... attribute to P, in both environments, representations of crackdows (thin, dark, marks that could be either shadows or cracks). (p. 208)

This is a particular instantiation of the general idea that we should attribute representations of "some neutral type ... that is satisfied by both [objects of the type that normally produce a particular kind of gray array on Earth] and [objects of the type that normally produce that kind of gray array on Twin Earth]" (pp. 202–3). My explicit use of disjunction is only intended to facilitate our description, as theorists, of the contents of P's perceptual experiences according to the strategy. Since we are dealing here with non-conceptual content, there is no suggestion that the subject is employing the concept of disjunction, nor that the visual system is manipulating an internal disjunction symbol.

We need to consider in some detail Segal's argument for adopting the strategy, and the consequences of adopting the strategy. But, before engaging upon that task, I shall briefly review three points that Segal regards as crucial to a proper understanding of Marr's theory of vision.

The first point is that: "*Each attribution of a representation requires a bottom-up account*, an account of how the representation was constructed" (1989, p. 194). The second point is that: "In general: *each attribution of a representation requires a top-down motivation*" (p. 195). The idea here is that it is illegitimate to attribute to a representation at one stage of visual processing a content that is more specific than can be used at any later stage. Recall that we think of visual processing in terms of inference. So, if we have already settled upon the representational content of the outputs of the visual processing system—as it might be,

that there is a cube four feet in front of the subject—then it would be gratuitous to assign to earlier stages of processing contents that are more specific than is required to license the inference to that conclusion. This way of constraining the contents of earlier stages in terms of their role in inference evidently obliges us to say something about the constraints on the attribution of content to the final stages of visual processing.

Segal's third point is that: "*attributions of representations are checked against behavioral evidence*" (p. 197). The idea here is that, if two different distal causes produce states such that the subject cannot, according to behavioural tests, discriminate between the two cases, then the two produced representational states should be given the same content. In particular, if two distal causes produce visual states that are of just the same intrinsic character and so have just the same behavioural consequences, then it would be illegitimate to credit the two states with different representational contents solely on the basis of the difference in distal causes. This third point may appear to be uncontroversial; nevertheless, it looms large in Segal's argument, as we shall shortly see.

*4.1 The argument for the strategy*

In order to review Segal's argument for the adoption of the strategy, we need to introduce some extra details into Burge's example of the shadows and the cracks (1986, p. 41; 1988a, p. 75). We suppose that there are particular circumstances W (let them be—as things actually are—abnormal and non-ideal circumstances: "W" for "wonky") in which an occurrence of C (a crack) is misperceived as O (a shadow). The circumstances W are to be contrasted with circumstances N ("N" for "normal"). We also suppose that the creature in the counterfactual situation is a duplicate R of the actual creature P. (For Segal (1989, p. 198), the counterfactual Visua* is a duplicate of the actual Visua.)

The circumstances W, which in the actual situation are abnormal and non-ideal, are normal in the counterfactual situation. According to Burge (1986), R sees occurrences of C (cracks) in circumstances W as C (as cracks), while P misperceives occurrences of C (cracks) in circumstances W as O (as shadows). Furthermore, P sees occurrences of O (shadows) in the actually normal circumstances N as O (as shadows); and we can suppose that in the counterfactual situation, if circumstances N were to obtain and instances of O (shadows) were to occur (perhaps counternomically), then R would misperceive those occurrences of O (shadows) as C (as cracks).

Now we come to Segal's argument. Imagine that P and R are in the same laboratory, and are subjected to two experiments. In Experiment 1, circumstances W obtain, and an occurrence of C (a crack) is presented to both subjects; in Experiment 2, circumstances N obtain and an occurrence of O (a shadow) is presented to both subjects. By Burge's lights, in Experiment 1, P is subject to an illusion while R is not; in Experiment 2, R is subject to an illusion while P is not. But, Segal argues, it is impossible to vindicate these differences between the intentional descriptions of P and R.

The reason is that "each attribution of a representation requires a top-down motivation", and in particular "attributions of representations are checked against behavioral evidence". (This is where that third point looms large.) For, in the laboratory, no behavioural differences between P and R will be discovered. Neither subject will be able to discriminate between an O (a shadow) in circumstances N and a C (a crack) in circumstances W; both subjects will be able to discriminate between an O (a shadow) and a C (a crack) if the two are presented in circumstances N. But without a top-down motivation—in this case, without a behavioural motivation in the form of a behavioural difference between P and R—the difference in content attributions cannot be justified.

Thus, according to Segal, a content-using scientific psychology must credit P's and R's experiences with the same content; so, the argument disallows Burge's anti-individualist example.

The common content that Segal offers is a disjunctive content. Applied to Experiment 1, his strategy is intended to have the consequence that the properly scientific specification of the content of the experience produced in P and in R is that there is an O-or-C (a crackdow). Since just the same kind of experience is produced in Experiment 2, the same content is attributed there.

This is the same as Matthews's suggestion that, when P perceives an O (a shadow) in the actually normal circumstances N, and when R perceives a C (a crack) in the counterfactually normal circumstances W, the two experiences have the same perceptual content. Both objects are perceived "as instances of an objective type that includes both O's and C's" (Matthews 1988, p. 83).

*4 2 Elaborating the strategy*

The basic idea behind the strategy is clear enough; but we can refine our appreciation of the strategy by asking what content is assigned to the experience produced in P by a C (a crack) in the actually normal circumstances N.

Recall that it is part of the set-up of Burge's example that in ideal circumstances P is quite able to discriminate a C (a crack) from an O (from a shadow); it is simply that P rarely sees Cs (cracks), and when he does it is in non-ideal circumstances (Burge 1986, p. 42). To be more explicit about this discrimination, I take it that in ideal circumstances P discriminates Cs (cracks) from Os (shadows) by distinguishing between them as such. This is to say that the properties of Cs (cracks) and Os (shadows) by which P discriminates between them are the very properties that make them Cs (cracks) and Os (shadows). For a surface feature to be a C (a crack) it must involve a certain pattern of change of curvature, while for a surface feature to be an O (a shadow) it must involve a certain pattern of change of shading on a flat surface.

In this case, there will be top-down (behavioural) motivation for distinguishing between the content of P's experience caused by a C (a crack) in circumstances N, and the content of P's experience caused by an O (a shadow) in circumstances N. This latter content is, according to the strategy, that there is an O-or-C (a crackdow). As things stand, the only very obvious candidate for the content of the former experience is that there is a C (a crack).

Suppose then that the strategy pronounces that, in circumstances N, P perceives a C (a crack) as a C (as a crack) and perceives an O (a shadow) as an O-or-C (as a crackdow). The problem is that these assignments of content utterly fail to make sense of P's behaviour on the discrimination task. *Ex hypothesi*, P can sort Cs (cracks) from Os (shadows) in circumstances N. This performance is unintelligible if the intentional description of P's experience has P perceiving the Cs (the cracks) as a subclass of the Os (the shadows). Yet to be a crack is surely to be a crackdow.

It might be replied that this is not unintelligible; for we can make sense of someone sorting dogs from (other) animals, for example. But we make sense of this performance by supposing that the sorter recognises the animals that are not dogs as animals that are not dogs. If we apply that idea to P, then we have P perceiving an O (a shadow) as an O-or-C but not a C (as a crackdow but not a crack); in short, as an O (as a shadow). But that is precisely the attribution of content that the strategy was designed to rule out.

There is another reply that might be offered on behalf of the strategy. It might be said that the experience produced in P by a C (a crack) in circumstances N is intrinsically different from the experience produced by an O (a shadow) in the same circumstances. So, it is perfectly intelligible that these two experiences could have different behavioural consequences; and there is no problem understanding how P could perform the discrimination task.

Now, it is correct that, if we are prepared to settle for a pure stimulus-response account of P's performance of the discrimination task, then all that is required is some intrinsic difference between the two experiences. But it would be a hollow victory for individualism if its locally supervenient content could not serve in intentional explanations of behaviour.

I think that it is clear that, if the strategy is to be applied to the example of the shadows and cracks as Burge sets it up, then the strategy must be elaborated. But equally it is clear what elaboration is needed. If experiences are to be awarded disjunctive contents, the disjuncts must advert to environmental circumstances. Thus, the content of P's (or R's) experience produced by a C (a crack) in circumstances W (or their experience produced by an O (a shadow) in circumstances N) is not simply that there is an O-or-C (a crackdow). Rather, the content is that there is an O (a shadow) in circumstances N or a C (a crack) in circumstances W.

Recall that (despite my omission of the hyphens) there is no suggestion here that the visual system is manipulating complex symbols. To the extent that the visual system's code is symbolic, the symbol that is tokened when P perceives a C (a crack) in circumstances W can be an unstructured symbol. But the extension of that symbol includes Cs (cracks) in circumstances W and Os (shadows) in circumstances N.

This elaboration of the strategy retains the feature that it provides a content that can be shared by the experiences produced in both subjects in both experiments. And it has the advantage that it allows us to make intentional sense of P's performance of the discrimination task.

*4.3 The consequences of the strategy*

The fix that I have just suggested for the strategy is, of course, intended as something of a poisoned pawn. I claim that, once we consider the application of the disjunctive content strategy to a variety of anti-individualist examples, we shall find it quite unattractive.

To construct an anti-individualist example in the context of Burge's second argument, for example, we alter the external environment in the counterfactual scenario so that experiences of the intrinsic type T normally produced by occurrences of property O (shadows) in the actual set-up are now normally produced by occurrences of property C (cracks). In principle, this feature of the counterfactual scenario could be produced either (a) by circumstances W being normal, or (b) by way of some more fundamental change in the behaviour of light, perhaps involving a change in the laws of optics. Thus far, we have followed option (a); but suppose that instead we opted for (b).

Suppose, in particular, that circumstances N are normal in the counterfactual situation as in the actual situation, but that, because of other changes, experiences of intrinsic type T are produced by occurrences of C (cracks) in circumstances N, but not by occurrences of C (cracks) in circumstances W. Furthermore, suppose that—as in Burge's own version of the example—there are no occurrences of O (shadows) in this scenario; and add that if there were occurrences of O (shadows)—perhaps counternomically—they would be invisible, or else would produce some quite other type of experience.

Now, what is the content of R's experiences of intrinsic type T in this wild counterfactual scenario? We have already found reason to say that the strategy awards a disjunctive content to P's actual experiences of the same intrinsic type: the content that there is an O (a shadow) in circumstances N or a C (a crack) in circumstances W. But, if R's counterfactual experiences are credited with that same content, then we fail to preserve a key feature of the strategy; namely, that experiences produced in normal circumstances are generally veridical. For R's experiences of intrinsic type T are never produced by an O (a shadow) under any circumstances, and are normally produced by a C (a crack) under circumstances N and not under circumstances W.

On the other hand, if we attribute to R's experiences the content that there is a C (a crack) in circumstances N, then we lose the most essential feature of the strategy; namely, that perceptual content is preserved across duplicates.

The friend of the strategy, having come this far, has only one obvious way forward. That is to expand the disjunction, and to complicate the disjuncts. For the content of R's counterfactual experiences—and so also for the content of P's actual experiences—a minimum of three disjuncts will be required; and each disjunct will have to mention objective properties—such as O and C—environmental circumstances—such as N and W—and other background conditions including even laws of nature.

But we cannot expect to stop at three disjuncts. For Burge can generate further potentially anti-individualist examples. One method is to take a counterfactual

scenario as home base, and to replay the considerations about objectivity and the possibility of error with respect to that scenario. Clearly, this method can be iterated. So, in general, there is no reason to think that the disjunctions can ever be completed.

The significance of this fact is not, of course, that the visual system is supposed to be burdened with the manipulation of indefinitely long formulae in its proprietary code. The significance concerns, rather, the content of P's actual experience produced in normal circumstances by a particular kind of shadow. According to Burge, the content of P's experience is that there is a shadow of a certain size and shape. The disjunctive content strategy—as elaborated to permit an intentional explanation of P's sorting behaviour—has led us to a very different content attribution. According to the strategy, P's experience has a content that would make it veridical if it were produced by a shadow of that certain size and shape under normal circumstances if the laws were as they actually are, or if it were produced by a similar sized crack under certain wonky circumstances if the laws were as they actually are, or if it were produced by a crack under the actually normal circumstances if the laws were thus-and-so, or ..., without end (cf. Egan 1991, p. 200).

## 5. *Evaluation of the strategy*

I claim that we can now assemble at least four reasons for rejecting the application of the disjunctive content strategy to examples of the type that Burge uses in his second argument.

### *5.1 The possibility of error*

First, the strategy bumps up against the very same "disjunction problem" that afflicts causal covariance theories of content (Dretske 1986).

At one point, Segal (1986, p. 202) supposes that Burge would argue against the strategy as follows. Given that a certain type of experience is, in a counterfactual scenario, normally produced by occurrences of C (by cracks), and given the success-oriented methodology of Marr's theory, the experience has the content that there is a C (a crack). And Segal points out that mere success-orientation does not vindicate the attribution of the content that there is a C (a crack) rather than the content that there is an O-or-C (a crackdow).

But Burge has more to say about the kind of strategy that Segal envisages. Matthews (1988, p. 83) floats the idea of disjunctive contents, and Burge replies using the premises of his second argument, and not just the assumption that experiences are normally veridical. What is crucial to this reply is the assumption of objectivity, and with it the possibility of error or misrepresentation. If the strategy is to establish individualism, then it must respond to every counterfactual example that Burge can construct; and Burge can construct a counterfactual example out of any possibility of what he calls "fundamental misperception":

> This is a misperception where the organism is unable in the context... to discriminate the thing as it perceives it from the thing as it actually is. (1988b, p. 97)

The strategy's introduction into the content of actual experiences of an extra disjunct in response to each new counterfactual example successively eliminates the possibilities of fundamental misperception; and in doing so it flies in the face of the assumption of objectivity.

*5 2 Marr's theory of vision*

Second, the application of the strategy to Burge's example is not supported by Marr's theory of vision.

Segal claims support from Marr because Marr's theory makes use of a notion of an *edge* that is somewhat similar to Segal's notion of a crackdow. According to Marr's theory, at a very early stage of visual processing—the primal sketch—edges are detected. At this stage, many very different objective features of surfaces are grouped together: an edge might be the boundary of a shadow on a smooth surface, or it might be the boundary of a surface. Later stages of visual processing sort out real surface discontinuities from variations in illumination, until ultimately the subject is presented in experience with information about objects with three dimensional shapes.

While the similarity between Marr's edges and Segal's crackdows is clear enough, it is not clearly relevant to the point of Burge's second argument, where the example of the shadows and cracks is actually introduced. There, Burge is primarily concerned with the content of visual experiences of three dimensional objects. So, when Burge talks about shadows and cracks he is concerned with a stage of visual processing far later than the stage at which edges are detected; roughly, he is concerned with the 3D model representation.

Invoking the requirement of "top-down motivation" (1989, p. 210), Segal notes that the representational contents attributed at these later stages of visual processing are answerable to the subject's ability to discriminate in experience between the different things in the world that are all grouped together as edges. He then applies the same requirement to cracks and shadows:

> Only if the subject could visually discriminate between them, under normal circumstances, would it be correct to attribute crack or shadow representations, rather than crackdow representations. (p. 210)

But this does not support the use of the strategy against Burge's second argument, since Burge is repeatedly explicit that his subject P is able, under normal or ideal conditions, to distinguish occurrences of C (cracks) from occurrences of O (shadows) (1986, p. 42; 1988a, p. 75). The example in Burge's second argument depends only upon the possibility of a particular occasion on which the subject is unable to make the discrimination:

> We do attribute visual representations of cracks or shadows, even where—*in a given instance*—the perceiver is fundamentally unable to discriminate the one from the other. (1988b, p. 98; my emphasis)

*5.3 Perceptual content and the explanation of behaviour*

Third, the strategy fails to deliver attributions of perceptual content that can figure in the intentional explanation of behaviour.

We have already seen one aspect of this problem in the fact that the strategy in its original form provides attributions of perceptual content that cannot figure in the explanation of sorting behaviour (§4.2). In essence, the problem is that, by aiming at local supervenience for the intentional content of experiences, the strategy conflates the intentional type of an experience with something that plausibly *is* locally supervenient, namely, the intrinsic phenomenal character of the experience.

There are two aspects of this conflation. On the one hand, every intrinsic difference comes to make a difference to content, since we can always construct a counterfactual scenario in which the intrinsic difference corresponds to a difference in distal causes under normal conditions. Consequently, if two experiences are intrinsically different, then the disjunctions used in specifying their contents will differ by at least one disjunct. However, we have no antecedently demonstrated use for a notion of perceptual content that cuts as finely as phenomenology. On the contrary, we have reason to think that perceptual experiences may differ intrinsically even though they are the same in point of content. Experiences may be sensationally different, though representationally the same (Peacocke 1983).

On the other hand, the perceptual content of an experience comes to be specified in terms of all the possible states of affairs that could be causes of such an experience. The indefinitely long disjunction that specifies all the possible states of affairs in which an experience would be veridical also specifies how the world seems to the experiencer to be. And the notion of how the world seems to be fits into a slot in the explanation of behaviour. Thus, for example, if it seems to a subject that there is a cube four feet in front of her, this seeming can be appealed to in the intentional explanation of the spatial behaviour of that subject: perhaps she moves forward about four feet and reaches out in such a way as to grasp a cube. But, if the strategy is applied to contents involving shape and distance properties, then it never really seems to a subject that there is a cube four feet in front of her. Rather it seems that there is a cube four feet in front of her in circumstances such-and-such if the laws are thus-and-so, or ..., or ..., without end. The objection to this is not that such a long disjunction never crosses the subject's mind. Rather, the objection is that such a seeming cannot contribute to the explanation of spatial behaviour unless the subject has information that allows her to rule out many of the disjuncts; and—given the strategy—such information cannot be gained from perceptual experience.

This problem for the strategy of attributing disjunctive contents is broadly similar to a problem that Burge launches against certain uses of definite descriptions in the specification of perceptual content, as in: the content of this experience is that there is a distal object of the type that normally causes this very type of experience. As Burge notes, this kind of content

> will not serve the needs of psychological explanation as actually practiced. For the descriptions of information are too inspecific to account

> for specific successes in solving problems in retrieving information about the actual, objective world. (1986, p. 38)

*5 4 Perceptual content and the theory of concepts*

Fourth, it is far from clear that the non-conceptual perceptual contents yielded by the strategy can be coherently related to the conceptualised contents of judgements that are based upon experience.

It is very plausible that a theory of the non-conceptual contents of experience should serve as input to a substantive theory of the conceptualised contents of thought and judgement. Thus Peacocke:

> We can consider the case of a possession condition for a relatively observational concept. It is plausible that such a possession condition will link mastery of the concept in question to the nonconceptual representational contents of the thinker's perceptual experience. (1989, p. 5)

The idea, very roughly, is that to grasp a concept that is at least partly observational—say, the concept of a cube—a subject must have a disposition to judge that an object is a cube if the object is perceptually presented *as* a cube. And—avoiding circularity—this latter notion is cashed out in terms of the *non-conceptual* content of the experience of the object being that there is a cube spatially related in such-and-such a way to the subject.

This theoretically desirable close mesh between concepts and non-conceptual content is not possible if there are no experiences in which an object is presented as a cube. And there will be no such experiences once the disjunctive content strategy has been applied.

For these four reasons, I reject the application of the disjunctive content strategy. This is not to deny that sometimes the content of a perceptual experience is most illuminatingly expressed in the theorist's language by a disjunction. What is being rejected is the employment of disjunctive contents in the service of the doctrine that—for scientific purposes—perceptual content must abstract from all environmental factors.

The strategy is a detailed development of the individualist's most promising line of attack, given that Burge's anti-individualist arguments are compelling against the conservative individualist. If that revisionary individualist strategy is rejected, then the anti-individualist arguments demand acceptance.

## *6. The error in the strategy*

It is one thing to reject the revisionary strategy, and another thing to identify the point at which the argument for the strategy goes wrong. The first is unsatisfying without the second.

I claim that the argument goes wrong in the particular way that it deploys the requirement that "each attribution of a representation requires a top-down [in this case, a behavioural] motivation" (Segal 1989, p. 195). That requirement is

deployed so as to rule out any intentional difference between duplicates tested in the same laboratory environment, even though their normal environments are quite different. In my view, this use of the requirement begs the question against the position that says that cognitive psychology treats information processing systems (modules) and whole creatures qua embedded in particular larger systems and ultimately in particular environments.

*6.1 An imaginary example*

Within psychology there are more specialised areas such as psycholinguistics or the psychology of vision. The psycholinguist, for example, studies the language system, holding fixed the presumption that visual and auditory processes deliver inputs to the language system, and that articulatory and other motor systems are available to receive outputs from the language system. Within psycholinguistics, there is further specialisation: to take just two examples, some psycholinguists focus upon the reading aloud of single words, others focus on the interpretation of pronouns in extended discourse.

Similarly, within the theory of vision, there can be further specialisation. So let us just imagine the psychological study of a particular subcomponent of the visual system—call it the *visex*. Since the visex is a component of an information processing system, the theory of the visex will make use of intentional descriptions. It might be, for example, that the input states of the visex contain information about the visible properties of distal objects—such as their depth—or about the properties of other internal states—such as binocular disparity.

Now imagine that there is, within the auditory system of some actual or imagined creature, a component that is physiologically the same (intrinsically) as the visex. Call this component the *audex*. In theories of auditory processing, some of the states of the audex are reckoned to be representations, and to have contents that concern, perhaps, properties of environmental objects such as sound waves, or properties of other internal states. Since visexes and audexes are intrinsically the same, the contents attributed to the states of visexes and audexes are not locally supervenient.

Consider now a particular visex V and audex A, and suppose that A is removed from its normal embedding in an auditory system and is plugged into the visex slot in a visual system. Various experiments are performed upon V and A and, of course, since "the twins are twins [they] will be the same in every testable respect" (Segal 1989, p. 205). Audex A is then returned to its normal location, and V is plugged into an auditory system. Once again, experiments reveal no behavioural difference between V and A.

*6 2 The lesson of the example*

From these experiments, it surely does not follow that the representational states of the visex should be credited with disjunctive contents—where one disjunct speaks of depth or disparity and the other disjunct speaks of properties of sound waves, for example— so as to allow our content-using psychology to generalise over these physiological twins. Rather, the psychological theory of the visex generalises over visexes qua components of the visual system. From the point of

view of the psychology of the visex, those accidental duplicates within some other creature's auditory system may be of no interest. If that is so, then in order to find generalisations that range over both visexes and audexes we need to descend to some lower-level, and more general, science.

Similarly, Burge's view is that the theory of the whole human visual system generalises over humans qua occupants of a particular kind of physical environment. The existence of accidental duplicates in very different environments may be of no interest to that theory; certainly, the mere possibility of such duplicates may be of no interest. The revisionary strategy begs the question against that view.

Of course, it *might* be that the discovery of the physiological match between visexes and audexes leads to further exploration of the visual and auditory systems in which they are respectively embedded, and to the discovery of theoretically important functional isomorphisms between the two. It might be that there develops a specialism within psychology that precisely studies the information processing in visexes and audexes. It might be that this specialism uses intentional descriptions that are somehow neutral between visual and auditory information. All this might be so: a specialism within psychology might use intentional descriptions of components in such a way as to abstract from *some* aspects of the embedding of those components in larger systems.

But, we can make three brief observations about these speculations. First, these intentional descriptions would still not be locally supervenient. Second, the possibility of these developments in psychological theory does not threaten the legitimacy of the original theory of the visex qua component of the visual system. Third, none of these speculations even begins to suggest that psychological theory ought to restrict itself to locally supervenient taxonomies of components.

According to Burge's view, what goes for the visex, the audex, and their respective embeddings in visual and auditory systems also goes for human beings, imagined duplicates, and their respective environmental embeddings.[1]

*Philosophy Department*
*Birkbeck College*
*Malet Street*
*London WC1E 7HX*

MARTIN DAVIES

## REFERENCES

Burge, T. 1979: "Individualism and the Mental". In *Midwest Studies in Philosophy Volume 4 Studies in Metaphysics*, P.A. French, T.E. Uehling and H.K. Wettstein, eds., Minneapolis: University of Minnesota Press, pp. 73–121.

——1986: "Individualism and Psychology". *Philosophical Review*, 95, pp. 3–45.

[1] The original impetus to write this paper came from conversations with Frank Jackson. Tyler Burge, Christopher Peacocke, and Gabriel Segal provided valuable comments on an early version. Further comments by Segal on a penultimate version substantially improved my understanding of his position, and occasioned eleventh hour changes to the paper. Work on the paper was begun at the Australian National University and continued at MIT. I am grateful to the Radcliffe Trust, the British Academy, and MIT for financial support.

——1988a: "Cartesian Error and the Objectivity of Perception". In *Contents of Thought*, R.H. Grimm and D.D. Merrill, eds., Tucson, AZ: University of Arizona Press, pp. 62–76.

——1988b: "Authoritative Self-Knowledge and Perceptual Individualism". In *Contents of Thought*, R.H. Grimm and D.D. Merrill, eds., Tucson, AZ: University of Arizona Press, pp. 86–98.

Davies, M. 1986: "Externality, Psychological Explanation, and Narrow Content". *Proceedings of the Aristotelian Society Supplementary Volume 60*, pp. 263–83.

——1992: "Perceptual Content and Local Supervenience". *Proceedings of the Aristotelian Society*, 92.

Dretske, F. 1986: "Misrepresentation". In *Belief Form, Content and Function*, R. Bogdan, ed., Oxford: Oxford University Press. Reprinted in *Mind and Cognition: A Reader*, W.G. Lycan, ed., Oxford: Basil Blackwell, pp. 129–43.

Egan, F. 1991: "Must Psychology Be Individualistic?". *Philosophical Review*, 100, pp. 179–203.

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

Fodor, J.A. 1986: "Individualism and Supervenience". *Proceedings of the Aristotelian Society Supplementary Volume 60*, pp. 235–62.

——1987: *Psychosemantics*. Cambridge, MA: MIT Press.

Jackson, F. & Pettit, P. 1988: "Functionalism and Broad Content". *Mind*, 97, pp. 381–400.

McGinn, C. 1982: "The Structure of Content". In *Thought and Object*, A. Woodfield, ed., Oxford: Oxford University Press, pp. 207–58.

——1989: *Mental Content*. Oxford: Basil Blackwell.

Matthews, R.J. 1988: "Comments [on Burge 1988a]". In *Contents of Thought*, R.H. Grimm and D.D. Merrill, eds., Tucson, AZ: University of Arizona Press, pp. 77–86.

Peacocke, C. 1983: *Sense and Content*. Oxford: Oxford University Press.

——1989: "What Are Concepts?". In *Midwest Studies in Philosophy Volume 14: Contemporary Perspectives in the Philosophy of Language II*, P.A. French, T.E. Uehling and H.K. Wettstein, eds., Notre Dame: University of Notre Dame Press, pp. 1–28.

Putnam, H. 1975: "The Meaning of 'Meaning'". In his *Philosophical Papers Volume 2 Mind, Language and Reality*, Cambridge: Cambridge University Press, pp. 215–71.

Segal, G. 1989: "Seeing What Is Not There". *Philosophical Review*, 98, pp. 189–214.

Smolensky, P. 1990: "Representation in Connectionist Networks". *Intellectica*, no. 9-10: *Modèles Connexionnistes*, D. Memmi and Y.M. Visetti, eds., pp. 127–65.