Proceedings of the British Academy, **83**, 49–77

PHILOSOPHICAL LECTURE

Unity and Objectivity

SUSAN HURLEY
Department of Politics and International Studies, University of Warwick, Coventry, CV4 7AL

1. Introduction and motivation

EXPERIMENTAL RESULTS which suggest that commissurotomy patients may support two separate centres of consciousness are familiar.[1] In this article I will do three things. First, I shall describe some quite recent research on commissurotomy patients that seems to me even more intriguing and striking than the types of result that are already familiar, for reasons that will emerge. Second, I shall develop some philosophical arguments about the unity of consciousness, which have a rather Kantian character (even though they cannot be found in Kant). These arguments have a Kantian character in that they try to show how making sense of the unity of consciousness requires an objective world, something outside of the contents of consciousness. Finally, I shall try to show how the neo-Kantian philosophical arguments can be related to the recent empirical studies of commissurotomy patients that I mentioned. The philosophical arguments may help us to understand why these cases are so puzzling, and may also bring us to identify and question some of the presuppositions about consciousness that we bring to these cases. Ultimately the course of my argument will turn against a traditional conception of consciousness.

Read at the British Academy 13 March 1993. © The British Academy 1994.

Rather than substantially altering the text in order to respond to the commentaries, I have added a few longish notes. Notes marked with '*' have been added by way of partial response to points made by the commentators.

[1] Though the interpretation of these results is controversial. For one of many summaries and commentaries, see Charles E. Marks (1981).

The background problem that brings the recent empirical work and the traditional philosophical arguments together is this. It seems necessary to distinguish in general between mental states that are *together* within one consciousness and mental states that are not thus together but are in *separate* consciousnesses. How should we understand this difference — the difference between the *togetherness or unity* of some mental states occupying a given stretch of time (such as my seeing your face and my hearing my own voice now while I am talking to you) and the *separateness* of other mental states occupying that same stretch of time (such as my seeing your face and your hearing my voice)? How should we distinguish between communication, coordination and duplication of information *between* separate units of consciousness, on the one hand, and activity *within* one such unit, on the other hand?[2]

There are some easy cases, of course. But can we find a principled account of the unity vs. separateness of consciousness that applies not just to the easy, normal cases, but to the hard, pathological cases as well? Perhaps scepticism about the unity of consciousness is correct: perhaps there are just various different differences between the cases, no right answers in some cases, and no unified phenomenon of 'the unity of consciousness'.

The sceptical view about the unity of consciousness may be correct; I do not claim that it isn't correct. Indeed, my arguments here might be regarded as one way of supporting such a view. But I think it would be surprising to many people if it were correct. Many philosophers assume that there is such a thing as the unity of consciousness, even if they are very sceptical about other aspects of a traditional view of the mind. For example, Derek Parfit, who is sceptical about personal identity over time, claimed in *Reasons and Persons* that a mind is divided when '... there is no single state of awareness of both of [two] ... sets of experiences'. He commented:

> It may be objected that these claims do not explain but only redescribe the

[2] Notice that the issue here is one about the unity of consciousness at a given time, not about personal identity over time. By 'at a given time' I do not imply that the micro-timing of consciousness makes sense in the way criticized by Daniel Dennett (1991), chapters 5 and 6. All I need to assume is that some relatively brief period can be occupied by various conscious attitudes or experiences, so that the question of whether they are co-conscious or not arises. It need not even be determinate for all conscious attitudes or experiences whether they occupy the given period, so long as it is determinate for some, concerning which my question arises.

> unity of consciousness in each stream. In one sense, this is true. This unity does not need a deep explanation. It is simply a fact that several experiences can be co-conscious, or be the objects of a single state of awareness. (Parfit, 1984, p. 250)

He held there that the unity of consciousness could not in his problem cases be explained in terms of ascription to a single subject or person, but also (and more importantly, from the point of view of my concerns here) that it needed no explanation. Someone might agree with him that it cannot be explained by reference to the identity of a person or subject, but still think it needs some other explanation. So, despite his scepticism about personal identity over time, there is a sense in which Parfit there took the unity of consciousness at a time for granted. Whether or not this is still Parfit's view, it is nevertheless representative of a certain tendency.[3] (An interesting question about Parfit's position is whether it *needs* to take the unity of consciousness for granted in order to develop his scepticism about personal identity and his preferred viewed in terms of survival and what matters.)

I make these comments partly in order to justify my assumption that it is worth investigating the notion of the unity of consciousness, even if our conclusion may turn out to be a sceptical one. Furthermore, it may be worth investigating because it may turn out that other sceptical views depend on taking the unity of consciousness itself for granted. If the unity of consciousness cannot ultimately be made sense of, or can be but only in certain ways — for example, by depending on some notion of an objective world, or on some normative notion of personal coherence — then other sceptical views that take the unity of consciousness for granted may themselves be undermined.

2. Some recent commissurotomy research

I'll now briefly describe the intriguing recent work on commissurotomy patients that I mentioned, which the philosophical arguments will then lead us back to. This work was done by Justine Sergent in Montreal, and a report of it can be found in her article 'Furtive incursions into bicameral

[3] Robert Fogelin, by the way, makes a related point about Hume. Fogelin claims that Hume was not likely to be worried about how we should assign two qualitatively identical perceptions to different minds or bundles, and remarks: 'He thought that he had immediate access to his own ideas and he simply took it for granted that these ideas were *his*' (Fogelin, 1985, p. 104).

minds' (Sergent, 1990). I assume familiarity with the method of lateralizing inputs to commissurotomy patients through fixation, which allows the experimenter to project different information to each half-brain. Sergent's series of experiments projected pairs of numerals simultaneously, one numeral to each half-brain, and asked patients to compare them. Either hand was allowed to reply by pushing a lever to indicate the correct response in a given task. The experimental conditions made it impossible for the hemispheres to cross cue — that is, to transfer information by external means, such as facial expressions or movements.

There are three relevant experiments for our purposes. In the first one, the task was to compare the two numbers and say whether they were the *same or different*. Whichever hand they used, the patients were *no better than chance* at this same/different comparison. This is what you would expect of a commissurotomy patient based on already well-known results. Suppose a '6' is projected to one hemisphere and a '7' to the other; since supposedly neither hemisphere has access to the other's information, how could either respond accurately to the same/different question?

But now consider the second experiment. Now the patients are asked not to say whether the numbers are the same or different, but which one is *higher* (in the sense of *greater*): they are supposed to press the lever on the side of the higher number, using either hand. Surprisingly, they find this higher/lower comparison 'very easy', show 'no hesitation' (p. 547), and are highly accurate. How, one wants to ask, could they make the higher/lower comparison, when they cannot make the same/different comparison? This seems especially puzzling, since the fact that the number 7, for example, is higher than some other number entails that they are not the same number.

In yet a third experiment, the same patients *were* able to indicate, with accuracy considerably above chance though not as high as in the second experiment, whether a pair of numbers they were comparing quantitatively were *equal* as well as whether one was higher than the other; but they were still unable, when the first experiment was rerun an hour later, to decide whether numbers were the *same or different*. That is, in the context of quantitative 'higher' or 'lower' judgements with no mention of the terms 'same' or 'different', judgements of equality could also be made; but in a context in which the numbers had to be identified as the same or different, judgements of sameness could not be made.

One natural interpretation of what is going on, which appeals to

Sergent herself, is that information about comparative quantity is somehow shared between the hemispheres, but not information about the identity of numbers, or the shape or name of numerals. Thus she suggests that the disunity affects certain categories of information but not others, even though there may be inferential connections between them. But if this is the case, then the question of how we should think of these patients isn't just the interpretative question that earlier commissurotomy results raised: do they have one consciousness, or two? Her results raise the further interpretative question as well: can they have a partially unified, partially disunified consciousness? This latter possibility, as Thomas Nagel pointed out in his article on 'Brain bisection and the unity of consciousness', is inconsistent with the way we normally think about the unity of consciousness (Nagel, 1979). Because they raise this further possibility of only partial disunity dramatically, Sergent's cases seem especially puzzling.[4]

*[4] Sergent herself does not make a claim about the structure of consciousness; she does not claim that her patients are conscious of both half fields at once, or that their consciousness displays a partially unified structure. Her interpretative suggestions are primarily at the level of information rather than consciousness. Nevertheless, her results are obtained in the context of a tradition of commissurotomy work that has raised the questions: one consciousness or two? and that has given rise to interpretations of commissurotomy patients as supporting two separate centres of consciousness (as well as interpretations according to which information in one hemisphere is unconscious, or absent from consciousness). (See for example Roger W. Sperry, 1990.) In this context, her suggestions about partial unity and partial disunity at the informational level naturally raise the further question about the possibility of partial disunity at the level of consciousness.

Later in the text I pursue this question in terms of a hypothetical case inspired by her results. It is important to recognize that the passage Marcel quotes from me in section 1 of his commentary is describing my own thought experiment, not any of her actual cases. Possible interpretations of that hypothetical case include interpretations in terms of (a) unconscious information, as well as in terms of (b) separate centres of consciousness with duplication of certain contents and in terms of (c) partial disunity or consciousness. I refer in the text to issue between these possible interpretations telegraphically as the separate-absent-or-nontransitive issue about consciousness. My primary concern here is with the issue between the separate and 'nontransitive' interpretations, but I admit the point that the 'absent from consciousness' (or unconscious information) interpretation is possible as well, and that my remarks about ease of response and there being no need to prompt guessing do not settle the matter. I do not deny that the 'separate-or-absent?' issue about commissurotomy patients is a difficult and important one, or that it is raised by Sergent's results. I am here indebted to Tony Marcel's commentary, which has pressed me to clarify my assumptions about the relevance of Sergent's work to the philosophical problems I consider.

An interesting thought experiment is this: what would happen if Sergent's patients were asked at the same time to do both the same/different and the higher/lower tasks? Sergent is careful to keep the two types of tasks separate. She takes care not to use the words 'same' and 'different' in giving instructions for the higher/lower/equal task and then interposes a mental rotation experiment before returning to the same/different task, in order to deter

3. Philosophical background: the unity of consciousness and the unities of concept and object

Having seen that some recent experimental work raises the possibility of interpretation in terms of partial unity of consciousness, let's now move back to the philosophical arguments. The two will connect up at the end. The argument I want to pursue is a neo-Kantian argument, although for various reasons it couldn't have been found in Kant himself. What does this mean? It is neo-Kantian in that it tries, as Kant did in the Transcendental Deduction of the *Critique of Pure Reason*, to argue from the very possibility of subjective experience or consciousness, and in particular from the need to make sense of the unity of consciousness, to certain objective requirements. Many philosophers hold that Kant's own arguments here failed, interesting as they are; and some philosophers have offered other arguments in their

---

subjects from resorting to the higher/lower/equal instructions in place of the same/different instructions (Sergent, 1990, pp. 548–9). However, consider a new task in which the same stimuli are presented, but subjects are asked both to press the button on the side of the higher number, or the middle button if the numbers are equal, and also to report verbally whether the numbers are the same or different. Would the dissociation persist, or would subjects learn to make consistent responses? If the latter, would a kind of adaptive unification of consciousness have been achieved? I discuss these matters further in *The Reappearing Self* (work in progress).

Colwyn Trevarthen (1984, p. 333) also reports results of tests on a commissurotomy patient that suggest unity at one level despite disunity at another. Trevarthen's patient was trying to do a left hand task relating to an object in the right visual field while fixating a central point. However, the patient '... could not respond and described the shape as vanishing immediately the movement began ... Here the patient's testimony was that the image of the object, initially seen, was blotted out of awareness in the left hemisphere the instant a movement, initiated by the right hemisphere, had started ... While the subject kept in readiness to respond, intently fixating ... [the object] was invisible, even when moved, until the instant it crossed the vertical meridian. Appearance of the corner of the card past the midline in the left visual field triggered a forceful and rapid response with the waiting left hand.'

Trevarthen (1974, pp. 247, 257) claims that the '... two-way links of hemispheres with the brain stem, as well as the links with the body itself, make complete surgical duplication of consciousness in man an impossibility' and that his findings, such as in the case in which the left hand intention seems to cause the right visual field object to disappear, '... require interpretation in terms of a representation of the visual field in the undivided parts of the brain that are in functional communication with both hemispheres at once ...' I discuss Trevarthen's findings and views, as well as Sergent's, at some length in *The Reappearing Self*. While Trevarthen's explicit claims come closer to raising the issues about partial unity of consciousness that I wish to address here than do Sergent's, his results are even more complex to interpret (for reasons related to the input assumption, which I touch on at the end of this paper).

place. The argument I'm going to pursue, an argument from the unity of consciousness to objectivity, comes into this category. Before developing the particular neo-Kantian argument I have in mind, however, some general philosophical background may be helpful.

We can distinguish three kinds of unity. The first, *unity of type* or *conceptual unity*, is the kind of unity involved when various objects all share some one attribute: fall under some one concept, are instances of some common property, are tokens of some common type. The second kind of unity, the *unity of objects*, is the kind of unity involved when various attributes all attach to the same object. The third kind of unity, the *unity of consciousness* at a time, is the kind of unity involved when various contents of consciousness at a time, including both experiential events and propositional attitudes, are all associated with one centre of consciousness, that is, are co-conscious.

Now it is easy and natural to assume that these three kinds of unity are at least partly independent of one another, and, in particular, that the unity of consciousness is independent of conceptual unity and the unity of objects and vice versa. However, Kant apparently wanted to deny this seemingly natural independence view by in effect equating the unity of consciousness with the combination of the other two forms of unity. He seems to see connections in every direction between the three unities, of consciousness, of concepts, and of objects. Not only does the unity of concepts require the unity of consciousness and vice versa, but the latter in turn require the concept of the object, and indeed experience of objects requires the unity of consciousness. In particular, he claims that there is an intimate two-way connection between our way of making sense of the identities of objects and our way of making sense of the identities of subjects.[5] He doesn't really seem to regard the three unities as distinct at all.

However, as various commentators have pointed out, Kant never made the full basis of his equation of the unities clear.[6] In particular, commentators have found arguments lacking or inadequate from the unity of consciousness to either something objective or the concept of objects, and in the other direction, from the concept of objects to the

[5] See and compare Quassim Cassam (1989) pp. 86–7 and *passim*; Quassim Cassam (forthcoming); Ralph Walker (1978) p. 76; Henry E. Allison (1983) pp. 144–5; Jonathan Bennett (1966) pp. 129, 131ff; Hubert Schwyzer (1990) pp. 90–98; C. Thomas Powell (1990) pp. 62–4.

[6] See, for example, Richard Rorty (1970) pp. 238–41; Quassim Cassam (1987) p. 375; see also Cassam (1989) p. 93; Walker (1978) pp. 81–2; Cassam (forthcoming); Allison (1983) p. 146.

unity of consciousness. And in both directions, philosophers have offered arguments to fill the gaps. For purposes of categorizing neo-Kantian arguments in this area, we can distinguish (at least) three types of argument. There are, first, arguments that the unity of consciousness requires objectivity, something outside of the contents of conscious experience.[7] Second, there are arguments that the unity of consciousness requires the *concept* of something objective, or objective content.[8] (Note that if arguments of the first type are successful, then the unity of consciousness requires something objective, not merely the concept of something objective.) Third, there are arguments that the concept of something objective requires the unity of consciousness.[9] While many interesting issues are raised by arguments of the second and third types, my concern here is with an argument of the first type, from unity to objectivity.

4. The 'just more content' argument from unity to objectivity

I want to begin by considering a version of an argument from the unity of consciousness to objectivity that derives from a discussion by Bernard Williams of Descartes' *cogito*. Descartes tried to argue himself out of extreme sceptical doubt by moving from his own doubt itself to his existence: I think, therefore, I am. A well-known objection to the premise 'I think' was made by Lichtenberg. He objected that Descartes had no right to the 'I think' premise of the *cogito*, but only to the

[7] See and compare P. F. Strawson (1959) chapter 3; Bernard Williams (1978) pp. 95–101. This latter text by Williams is the starting point for my discussion in section 4 below.

Note also Margaret Wilson's description of Descartes' position: he '... provides no account at all of what individuates souls, or distinguishes one from another'. To make his position at all plausible, she suggests, he would have to introduce a principle that would not be easy to justify: '[h]e would have to hold that mental substances may after all be identified and distinguished according to the psychological or phenomenological features they "exhibit" '. You and I can claim that '... we are different thinking things on the grounds that we have (or think we have) different and incompatible thoughts, memories, attitudes, and so forth'. This principle would provide a Cartesian version of the view that the neo-Kantian 'just more content' argument developed below cuts against. See Margaret D. Wilson (1987) p. 199.

[8] Strawson (1966) pp. 100–101. See Cassam (forthcoming); Walker (1978) pp. 117–19; Rorty (1970) p. 218.

[9] See and compare, for example, Cassam (1989) pp. 89–92ff, 95–6; Cassam (forthcoming); Allison (1983) p. 146; Strawson (1966) p. 103.

impersonal premise 'there is thinking going on'. 'I think' implicitly claims something more substantial about what is objectively the case, in referring to a thinker, than is warranted just on the basis of the subjective materials, the existence and content of conscious thought itself, which Descartes allowed himself. There is some affinity between Lichtenberg's objection and Hume's views about the self.

Williams responds to this objection of Lichtenberg's, and explains why it doesn't work. He points out that, given any two thoughts that *p* and that *q*, it must be possible for them to be *separate*, so that there is no thought that *p* and *q*, as well as *together or co-conscious*, so that there is a thought that *p* and *q*. If the two thoughts belong to separate consciousnesses, for example, there is no difficulty in understanding how they can be contradictory. It must be possible for there to be more than one 'thought-world', to use Williams' phrase; in one it may be thought that *p*, in another it may be thought that not *p*. Given that separateness or disunity of consciousness must be possible (as it is with different people), then we need some way of understanding this possibility in general, even when *p* and *q* are not contradictory. But the impersonal Lichtenbergian 'there is thinking going on' doesn't provide what is needed. Basically, this conception of the subjective is too weak to allow us to make sense of the togetherness or separateness of mental states. But making sense of this seems to be an essential part of making sense of the subjective realm: the unity/separateness issue arises and needs some resolution. Therefore, this Lichtenbergian/Humean conception of the subjective must be too weak.

So, I want to pursue this general issue about the unity vs. separateness of consciousness. Here is another way of putting the issue. If it is thought at some time *t* that *p* and it is also thought at that time that *q*, what determines whether it is also thought at that time that (*p* and *q*)? That is, what is it that determines whether two conscious states are unified, together, co-conscious, or whether they are separate? Whatever this is, let's make a placeholder index for it, '*i*', so we can say:

> if it is thought that *p* in *i* at *t* and thought that *q* in *i* at *t*, then it is thought that *p* and *q* in *i* at *t*.

We can call this the *agglomeration principle*. But it means no more than this: the issue about unity arises, and something settles it; call that something, whatever it is, '*i*'. Contents of consciousness agglomerate just when the indices of agglomeration, the '*i*'s, match. The agglomeration principle in itself says nothing about what '*i*' refers to; it is simply

a placeholder index for whatever governs whether the two thoughts agglomerate or not. Agglomerated contents exist just when the index is the same for *p* and for *q*.[10]

Consider then what happens if it is not thought that *p* and *q* in *i*, say because '*p*' and '*q*' contradict one another. Then it can't be thought that *p* in *i* and also thought that *q* in *i*; one of these contents must be absent from *i*, even though it may be present somewhere else, say in *j*. This is basically the formal pattern of our thinking when we attribute disunity to commissurotomy patients: for example, the patient says he sees just one point of light and it is red, and also indicates that he sees just one point of light and it is green. It is not possible for there to be an awareness of just one point of light that is both red and green. So we say that consciousness of the green light is separate from, not co-conscious with, consciousness of the red light (if conscious at all). The indexed agglomeration principle merely represents this commonsense way of interpreting commissurotomy patients formally. We distinguish between indices such as '*i*' and '*j*', which means no agglomerated content is entailed, because there is a normative prohibition on certain agglomerated contents, such as awareness of one point of light that is both red and green.

I am not, by the way, claiming that a noncontradiction constraint necessarily applies to all contents of consciousness, only that when it

*[10] Perhaps the agglomeration principle is stronger than it needs to be to raise the issues about strong unity that I wish to raise below under its head; perhaps it goes further than commitment to a transitive relation of co-consciousness goes. (See and compare Roderick Chisholm, 1981.) For example, as Michael Lockwood points out in his commentary, iteration of the principle of agglomeration gives rise to thoughts that *p* and *q* and (*p* and *q*) and [*p* and *q* and (*p* and *q*)] and . . . , *ad infinitum*, and this may be regarded as implausible. But no similar iteration seems to afflict mere talk of co-consciousness. From the co-consciousness of contents that *p* and that *q*, it does not follow that the further content that *p* and *q* is also co-consciousness with the contents that *p* and that *q*.

Now similar problems of iteration may also arise for certain conceptions of the logic of *knows that*. . . . But such problems may be less serious for the heavily normative concept of knowledge, where they may be regarded as merely reflecting a certain idealization of knowledge, than they are for the less heavily normative concept of consciousness, where no such idealization is evidently appropriate.

Perhaps the iteration problem could be blocked by adding a reflective component to the formulation of agglomeration to keep track of levels, plus something like a type-rule preventing cross-level contents. Or perhaps it could be blocked by invoking set theory directly, so that we spoke of consciousness of sets of experiences rather than of conjunctive contents, and a type-rule again ruled out cross-level contents.

On these matters I am grateful for comments by Michael Martin and Christopher Peacocke as well as to Michael Lockwood.

does apply it can determine that certain contents are separate. This is reflected formally by contraposing the indexed agglomeration principle in the way I just explained.

Someone might wonder why the agglomeration principle needs to be indexed. Consider the results of denying this. It may seem that solipsism would necessarily result: a state of affairs in which, as Williams puts it, '... there is ... only one ... point of view: events either happen for it, or they do not happen, and there is no way of conceiving of such events happening, but happening (so to speak) elsewhere.' We can spell this out as follows. If agglomeration held with no restriction, no indexation, then the contents of any existing mental states would be united in the one and only thought-world. Applied to conscious states, this means that for there to be consciousness of any two things at all, there would also have to be consciousness of both of them — they would *have* to be co-conscious, not separately conscious. If they were not co-conscious, one of them could not be present to consciousness at all. Now if we consider solipsism to be absurd, then the threat of solipsism may answer the question why agglomeration needs to be indexed.

However, what I have just said assumes the only options are: indexed agglomeration, or agglomeration without indexing, unrestricted. But another option might be to deny that agglomeration necessarily applies at all. I am going to return to this possibility later. For now I shall assume indexed agglomeration does hold.

If it holds, then the index '*i*' is just a placeholder for whatever it is that determines the unity or separateness of consciousness. Now we saw in one case that a normative prohibition could do this: there can't be awareness of one point of light that is both red and green; this involves an inconsistency. This is reflected formally by our distinguishing two indices, say '*i*' and '*j*', and recognizing separateness of consciousness. But the more general question still arises: what is the nature of the needed indices of agglomeration, what do they refer to? Perhaps in some cases they reflect conceptual norms, but it is doubtful this will always do the needed work. This is because it seems possible for two contents to be in separate consciousnesses even though they are perfectly consistent with one another logically or normatively. They may even be duplicates of one another: different tokens of the same type. So normative coherence does not entail co-consciousness.

The basic issue here is whether whatever does the needed work of determining the unity or separateness of consciousness, to which the

indices of agglomeration refer, can be based solely on *subjectively available resources*, that is, on resources internal to the contents of consciousness as traditionally conceived. To go back to Lichetenberg's weakening of the 'I think': can we get anything rich enough to determine the separateness or unity of consciousness if we restrict ourselves, with Descartes, to the subjective viewpoint? Or, will it be necessary to move outside the contents of consciousness, to something objective, in order to make sense of the unity or separateness of consciousness? This would be a neo-Kantian argument in that it would be an argument from the unity of consciousness to objectivity. For example, perhaps the indices of agglomeration reflect the objective identities of persons — the difference between you and me — or differences between bodies, or differences in spatio-temporal location, or neurobiological differences in the brain.

Let us first consider what can be done to determine the separateness or unity issue using just resources internal to content. I have already commented on the way in which normative relations among contents can determine relations between indices: contradictory contents can prompt attributions of contents to separate thought-worlds. But this normative approach to the problem will not yield separate thought-worlds in cases where no inconsistency or other normative failing is involved. We cannot make do with the principle that consistent contents all belong to the same thought-world; coherent contents may nevertheless belong to separate thought-worlds. Indeed, as I have already noted, conscious states of the very same content-*type* may be duplicated as distinct *tokens* in separate thought-worlds. A normative requirement of coherence among contents in any one thought-world will not help us to identify these separate thought-worlds. So this way of appealing to relations among contents to determine unity or separateness isn't sufficient, though it certainly has an important role to play.

Williams considers another possibility, namely, that adding to the contents whose separateness is in question a context-dependent term, such as 'here' or even 'I', might determine separateness or unity. Perhaps Descartes was entitled to more than the Lichtenbergian 'It is thought: *p*'. Perhaps he was entitled to 'It is thought: *I* am thinking that *p*'. This is still weaker than what Descartes wanted, namely, something of the form: 'I am thinking: *p*'. It is weaker because, while the content of the thought is personal (that **I am thinking that** *p*), the state of affairs involved in its being thought is still impersonal (**'It is thought**

that: I am thinking that *p*', not 'I think that: *p*'). So, even if Descartes might not be entitled to 'I think', he might nevertheless be entitled to more than 'there is thinking going on', namely, to 'there is thinking that: I am thinking'. This intermediate possibility provides in effect a weakened version of the Lichtenbergian objection. It enriches the content of the thought by making it personal, without committing itself to the existence of the thinker. Will this intermediate possibility be rich enough to provide a way of understanding the unity of consciousness? If so, we would have got round the neo-Kantian argument that something objective, something outside the content of consciousness, is needed.

But Williams explains that the weakened version does not, by enriching the *contents* of the thoughts, overcome the problem about separateness that faced the original objection. In my terms, such enriched contents do not provide a way of indexing the agglomeration principle, or of determining issues of separateness and unity. The truth that *it is thought: 'I am thinking that p'*, plus the truth that *it is thought: 'I am thinking that q'* do not together determine that it is also thought '*p* and *q*' or 'I am thinking that *p* and *q*', unless the same person is thinking in each case. But the latter is an objective matter. In the absence of objective context, first-person or 'I' contents do not determine the unity of consciousness; they are *just more contents* for which the unity issue arises again. For this reason, I think it's appropriate to dub this neo-Kantian argument the 'just more content' argument from unity to objectivity. The point is that *more* than 'just more content', even if it is 'I'-content, is needed to make sense of unity and separateness; something outside of content, something objective, is needed.

Note that it would not help to use the name of a person in the contents as well as 'I', since a version of the point just made will still apply. One thought-world may include the thought that I, who am the one and only Susan Hurley, am thinking that *p*, while another thought-world includes the thought that I, who am the one and only Susan Hurley, am thinking that *q*. There may still be no thought-world that includes the thought that I, the one and only Susan Hurley, am thinking that *p* and *q*, or even any thought-world including the thought that *p* and *q*. There will be if it is the same Susan Hurley who has both thoughts at the same time, but this again is an objective matter, not something entailed by the existence of thoughts with the given contents at the same time. (Each of two Susan Hurleys might falsely believe she was uniquely named.) Even if contradictory contents must be

separate, nothing we put into the contents whose separateness is in question will *in the general case* determine the unity or separateness issue or effectively index the agglomeration principle.

What apparently will do the trick of effectively indexing agglomeration is reference to something objective, such as a person or subject in the statements of the occurrence of the thoughts. 'Susan Hurley thinks that *p*' and 'Susan Hurley thinks that *q*' warrant 'Susan Hurley thinks that *p* and *q*', given that the same person or subject is thinking both token thoughts at the same time. Now it is possible, as Williams allows, that something objective but other than a person or subject will provide the needed index to govern agglomeration. But the point at present is not whether persons or subjects in particular are needed, but rather that what is needed will not be got from the contents of consciousness, from subjectively accessible materials, alone. In this sense, the contents of consciousness are not autonomous. In order to make sense of consciousness or the subjective realm itself, we must appeal to something beyond the contents of consciousness, to something objective. This is because it must in general be possible to determine whether contents of consciousness are separate or united and something objective is needed to do this. What that something else is remains open.

5. A naive objection to the just-more-content argument, and weak vs strong unity of consciousness

Let us now go back to a possibility I promised to return to: maybe agglomeration doesn't hold in either form, either indexed or unrestricted. Perhaps both forms are too strong. It is the assumption that agglomeration needed indexing, that the unity or separateness of thought-worlds needed determining, that leads into the argument that something objective is needed to do this work. But giving up indexed agglomeration need not mean falling back on unrestricted agglomeration and solipsism, so that there is just one big thought-world. Maybe we don't need agglomeration at all. In this way we can work ourselves back to ask an apparently naive question: why is there a problem about what determines separateness or unity of contents at all? Why must there be some general principle — either indexed, or unindexed agglomeration, for example — that determines separateness or unity?

Let us admit for the sake of argument (even if we later bring it

into doubt) that for any two particular states of consciousness — any two particular thoughts or experiences, for example — it must be determinate whether they are co-conscious or not. Maybe this much determinacy is necessary in order to make sense of the notion of consciousness. Furthermore, we can agree that co-consciousness is normatively constrained in the weak sense that contradictory or uncoordinated contents of certain types of state, such as belief, cannot be co-conscious, must be separate. But none of this entails agglomeration, either indexed or unrestricted. Instead, each particular case of co-consciousness, the objection goes, can be determined from the perspective of consciousness itself: if there is a consciousness that *p* and there is a consciousness that *q*, and we want to know whether there is a consciousness that *p* and *q*, there is no general or normative principle that generates the answer. Rather, it may be an open question: there may be a consciousness that *p* and *q* or there may not be; the perspective of consciousness must be consulted directly on this conjunctive content itself. There may be no independent principle that determines the unity or separateness of the contents of conscious states. The just-more-content argument goes wrong, according to this objection, in assuming that the perspective of consciousness has to generate an answer to the question about whether there is consciousness that '*p* and *q*' on the basis of something else. But it doesn't; it can simply help itself directly to awareness of the conjoined content '*p* and *q*', or absence of any such awareness. This is all that is needed for philosophical purposes. Even if co-consciousness depends on objective neurobiological facts, we don't need to consult neurobiology in order to determine co-consciousness.

What should we make of this naive objection to the neo-Kantian just-more-content argument? I am going to try to answer this question in terms of a distinction between strong unity and weak unity of consciousness. The consequence of the objection would be that we should give up the assumption that consciousness at a time displays strong unity, and admit the possibility of mere weak unity. The question then is: if we don't assume strong unity, will we be able to argue from unity to objectivity in the manner of the just-more-content argument? Perhaps not exactly. But I will try to show that even weak unity yields an argument to objectivity.

To explain the distinction between strong and weak unity of consciousness, I need first to point out the close relationship between the indexed agglomeration principle and another formal principle, that of

the transitivity of co-consciousness. These two principles both provide an initial characterization of what I will call the strong unity of consciousness. Here are instances of them:

Indexed Agglomeration:	**iCp & iCq → iC(p & q)**
Transitivity:	**pCCr & rCCq → pCCq**

What is the relationship between indexed agglomeration and transitivity? I shall first remind you of how indexed agglomeration works, then comment on transitivity, then on the relations between them. Keep in mind that all conscious states are assumed to occur at the same time; I am not concerned with transitivity over time but at a time.

We have already seen how the indexed agglomeration principle reflects our commonsense reasoning about commissurotomy patients. If the patient manifests consciousness that *p* and consciousness that *q*, but these are inconsistent, then we partition consciousness to reflect this inconsistency. In laboratory demonstrations of dissociation in commissurotomy patients, such inconsistencies are typically induced experimentally. For example, the patient is instructed to write down any and all numbers perceived; the right hand writes '*6*', and the left (probably with difficulty, even where language is possessed by both hemispheres) writes '*7*'. Or such contradictions are inferred from motor responses such as pointing or button pushing. Formally, we argue contrapositively: since we cannot make sense of consciousness that *p* and *q*, given the inconsistency, we attach different indices to consciousness that *p* and consciousness that *q*, so that agglomeration doesn't apply (and thus isn't violated). This formal move may be interpreted as partitioning the subject of consciousness: no one subject is conscious that *p* and *q*; instead there may be two subjects, Lefty and Righty, one conscious that *p*, the other conscious that *q*. Another interpretational option, however, is to recognize only one centre of consciousness; inconsistent content may then be unconscious rather than separately conscious. Some commissurotomy researchers think this is the right way to interpret these cases. At any rate, I take it that this latter assumption is the one we tend to make about blindsight and the large range of related cases in which the ability to use information implicitly seems to be dissociated from explicit awareness of it. So the question '*separate* consciousness or *absent* consciousness?' arises, driven by the agglomeration principle: these two options are two ways of avoiding a violation of agglomeration and of the strong view of the unity of consciousness it reflects. I shall now explain how the indexed agglomeration principl

in turn seems to involve another assumption about the unity of consciousness, namely, that the relation of co-consciousness between conscious states at a time must be transitive.

Let *x*, *y* and *z* range over things like particular experiences or particular propositional attitudes occurring at a given time, things we are conscious of or things we are conscious that. (I am not assuming that all propositional attitudes are conscious attitudes, but merely including any that are in this range.) These token experiences or attitudes, partially identified by the types of their contents, are the relata of the co-consciousness ('*CC*') relation. Then we tend to assume that co-consciousness is transitive:

(x)(y)(z) (xCCy & yCCz → xCCz).

For example, if a perception that the grass is wet and a perception that the sky is blue are co-conscious, and if the same perception that the sky is blue and a perception that the wind is strong are also co-conscious, then those perceptions that the grass is wet and that the wind is strong are also co-conscious. A unit of consciousness at a time, a unitary state of awareness, includes all the members of one group of transitively connected relata of the co-consciousness relation at that time. The index of agglomeration reflects such a unitary state of awareness: the index might refer, for example, to the subject of such a unitary state, or some other objective factor the transitively related states have in common.

To see that indexed agglomeration is very closely related to transitivity (for states occurring at the same time), notice that if the co-consciousness of states at a time were not transitive, it would not define a unitary state of awareness, some feature of which could serve to index agglomeration. If *x* were co-conscious with *y* and *y* with *z*, but *x* were not co-conscious with *z*, then we could not use the subject *i* of these states, for example, or something else that unites them, as a basis for agglomeration, could not say that since *i* is conscious of *x* and *i* is conscious of *z*, then *i* is conscious of (*x* and *z*). Why not? Because, since transitivity fails, *x* and *z* are not co-conscious. Without transitivity, even if the various contents all belong together in some sense, this will be weaker than what indexed agglomeration demands. Agglomeration builds bridges of co-consciousness between any two conscious states that share the relevant unifying factor, reflected by the index of agglomeration; but failure of transitivity allows for merely piecemeal unity of co-consciousness. If we allow this is possible, then to the separate-or-

absent question about abnormal cases we must add another option: nontransitive?[11]

So, it is this merely piecemeal co-consciousness that is appealed to by the naive objection to the just-more-content argument. Whatever does the job of the index of agglomeration in strongly unifying consciousness may not be open to the perspective of consciousness; but why do we need this anyway? All we need is to determine whether co-consciousness holds piecemeal, on a case-by-case basis, and this *can* be done from the perspective of consciousness. True, that gives us no guarantee that co-conscious will turn out to be transitive, but so what? It may not be in some cases. The unity of consciousness may be only weak or partial. Then agglomeration won't hold, but also the just-more-content argument to objectivity won't get started.

At this point we may be tempted to think that we can construct an index of agglomeration, and hence provide a way of understanding the strong unity of consciousness, out of a transitive series of co-consciousness relations, where each instance of co-consciousness is open to the perspective of consciousness. This might seem to be a way of getting around the neo-Kantian just-more-content argument on its own terms, without giving up strong unity, by resisting the claim that nothing in the content of consciousness can do the work of an index of agglomeration in determining the strong unity of consciousness. But this temptation involves a confusion, in particular, a confusion between individual instances or tokens of conscious states, and types of conscious states. Transitivity of co-consciousness applies to token states of consciousness, such as individual experiences or thoughts, not to the types of such

[11] There is more to be said about the precise logical relations between indexed agglomeration and transitivity. If a token conscious state can have more than one index, then we might have indexed agglomeration and nontransitivity. On the other hand, whatever necessitates transitivity, if it is necessary, might be what provides the one index of agglomeration of each state.

Assuming that each token conscious state had one index, however, agglomeration seems to be stronger than (to entail but not be entailed by) transitivity of co-consciousness for states at a time. Agglomerated content requires co-consciousness of the conjuncts agglomerated, even if co-consciousness doesn't require agglomerated content. As mentioned in an earlier note, talk of co-consciousness, even if constrained by transitivity, does not seem to give rise to the problem of iteration that agglomeration gives rise to. To the extent indexed agglomeration and transitivity of co-consciousness come apart, my characterization of strong unity must be refined. While my route here to the issue between strong and weak unity is via agglomeration, it may be that agglomeration raises difficulties not relevant to my present project. In that case, the basic characterization of strong unity should be in terms of transitivity of co-consciousness rather than agglomeration.

states, such as the painfulness of experiences. If *I* am conscious that *p* and *r* and *you* are conscious that *q* and *r*, we are both conscious of the content-type '*r*', but transitivity does not apply, because the token conscious states with contents of type '*r*' of which you and I are each conscious are not the same. What is open to the perspective of consciousness in this context is the type of the experience, but not its token identity. Token conscious states are only partially identified by the types of their contents. In general, then, even if there is consciousness that *p* and *r* and consciousness that *q* and *r*, the question of whether there is consciousness that *p* and *q* is not determined by simply applying transitivity. To know whether transitivity applies, we would have to know whether the conscious states with content '*r*' were the same token state as well as the same type of state, and this is not something that is open to the perspective of consciousness. If the states with content '*r*' were merely duplicates, not identical, then transitivity would not apply. Therefore, we cannot construct an index for agglomeration using transitivity of co-consciousness and relying merely on the contents of consciousness. Indeed, on the contrary, it seems that the identity or distinctness of token conscious states must in turn depend on something objective. As Gareth Evans wrote:

> A subject can gaze inwardly with all the intensity he can muster, and repeat to himself 'this pain', 'this pain', as he concentrates on his pain, but he will not thereby be able to know which pain is in question unless this provides him with a basis for identifying the pain with a pain conceived as an element in the objective order . . . (Evans, 1982, p. 253)

Suppose we follow the naive objection this far: we admit the idea of piecemeal co-consciousness, which is open directly to the perspective of consciousness, even though that provides no guarantee that co-consciousness will be transitive. Then we have given up the necessity of indexed agglomeration and the strong unity of consciousness, so the just-more-content argument from unity to objectivity can't get started: if nothing needs to play the strongly unifying role of the indices of agglomeration, we can't argue that something objective needs to play this role. But there is still a weaker sense in which consciousness might be unified, which allows failures of transitivity. The subject of such a *weakly unified* consciousness could suffer from merely local dissociations, merely partial disunity. A single instance of disunity would not then cut a swathe through the whole structure of consciousness as it does under the assumption of strong unity. Each token conscious state of that subject

at that time will be co-conscious with *some* other token conscious state of that subject at that time, rather than with *every* other. Whereas Parfit's fission cases get us to consider the possibility that consciousness over time may have a branching structure, we are now considering the possibility that consciousness at a time may have a branching structure. Then inconsistent statements or actions won't necessarily prompt us to partition consciousness globally or to suppress certain contents from consciousness, as they do under strong unity. There is now a third option: inconsistent statements or actions may reflect merely local disunity. The two-fold 'separate-or-absent?' issue becomes the three-fold 'separate-absent-or-nontransitive?' issue. As a result of giving up the assumption of strong unity, the power of normative constraints to determine unity or disunity is greatly reduced, rendered merely local.

Recall that the co-consciousness relation can be determinate in a piecemeal way — either hold or fail in each particular case — even if it is not transitive. Thus we could regard determinacy of co-consciousness to be an essential characteristic of consciousness, while allowing that transitivity may not be, so that indexed agglomeration and strong unity cannot be assumed as necessary. In this way we have now come back to our question: what are we to make of the naive objection to the just-more-content argument from unity to objectivity? It seems we can appeal merely to relations of co-consciousness, which are open, for content-types at least, to the perspective of consciousness. We cannot escape the requirement of objectivity by imposing transitivity on these relations if, as it seems, the application of transitivity to token states of consciousness itself requires objectivity. But instead we might simply *give up* transitivity as well as agglomeration, and appeal to the perspective of consciousness directly in each case to determine whether content types were co-conscious or not. Nothing objective is needed to determine that.

6. Nontransitivity vs. duplication: a reply to the naive objection

At this point I want you to recall Sergent's commissurotomy cases, which I described at the beginning. Recall that these patients could make higher/lower comparisons, but could not make same/different comparisons. This seemed to suggest to Sergent that disunity may be limited to certain types of information or ability. But that suggestion raises the

further possibility that the unity or disunity of consciousness may not be not all or nothing. Now it is one thing to interpret such a case in terms of limited disunity *of information*, and another thing to interpret such a case in terms of limited disunity *of consciousness*. We shouldn't lose sight of the possibility that some of the information in question is unconscious. The separate-or-absent issue is itself difficult and important, and I do not pretend to have done it justice.[12] But to the extent that it is ever prima facie warranted to interpret commissurotomy patients as supporting two separate centres of consciousness rather than in terms of one centre of consciousness plus some unconscious information, we can at least frame a hypothetical case which assumes that prima facie warrant and combines it with the further puzzling features of Sergent's cases. Such a hypothetical case would then raise the philosophical issue about possible weak disunity at the level of consciousness that I want to consider, even if it is not raised by Sergent's actual data.

I attempt to frame such a hypothetical case as follows. Suppose a '*6*', on the left hand side, is projected to the right hemisphere and a '*7*', on the right hand side, is projected to the left hemisphere. Assume that either side can report reliably, using the contralateral hand, the number projected to it, and that we are warranted in supposing that there is a conscious perception of a number as a *6* and that there is a conscious perception of a number as a *7*. But, since the patient is unable to indicate whether the digits are the same or different, presumably the perception of a number as a *6* is not co-conscious with the perception of a number as a *7*. Now the patient is able to indicate that the number on the right is higher than the number on the left with ease and no need to force a guess, as there is in blindsight. This observation may go some way to counter the suggestion that the comparative information must be implicit or unconscious, though it is not decisive. But for purposes of my thought experiment (as opposed to Sergent's actual experiment) I want to focus not on the issue about consciousness vs. unconscious information, but rather on the further interpretative issue about weak

*[12] Some of Marcel's work, discussed in his commentary, takes us further toward understanding how this issue should be addressed. Questions remain open as to how many years of practice at supposedly guessing might affect the distinction between guessing and conscious report. Marcel's work also raises further versions of the issue about possible partial disunity of consciousness that I am concerned with, though I am not sure I understand Marcel's distinction between access and synchronous co-reference (section 5). See also Michael Gazzaniga (1988) on commissurotomy patients who are able to understand language and speak from each half-brain. And see Sperry (1990); Sperry also discusses the difficulty of resolving the issue about partial unity vs. duplication that is my focus here.

or partial unity that arises on the assumption that the relevant information is conscious. So I presume that the patient also consciously perceives that the number on the right is higher that the number on the left. And since there is no indication of dissociation within either the right or the left hemisphere considered separately, presumably the perception of the number on the right as a *7* is co-conscious with a perception of it as higher, and the perception of the number on the left as a *6* is co-conscious with a perception of it as lower.

One possible interpretation of this hypothetical case is as involving weak unity and nontransitivity of co-consciousness: there is unity with respect to awareness of comparative quantity but disunity with respect to awareness of the identities of the digits. To spell this interpretation out: awareness of the number on the left as lower is co-conscious with awareness of the number on the left as a *6* and also co-conscious with awareness of the number of the right as higher, which in turn is co-conscious with awareness of the number on the right as a *7*, even though awareness of the number on the left as a *6* is not co-conscious with awareness of the number on the right as a *7*.

I have used the Sergent results to generate a thought-experiment that illustrates the possibility of weak unity, by combining her own suggestions about the possibility that disunity might be limited to certain types of information with the further assumption that the relevant information is conscious information. But there may well be other actual or possible cases that provide as good or better illustrations.[13] Perhaps the strongest urge to attribute weak unity may be associated with cases where the neuroanatomy is known to be normal or very close to normal: consider for example a patient who manifests only very limited dissociations after surgery that severs only a very small portion of his corpus callosum. In such a case, assuming that both of a pair of dissociated perceptions should be regarded as conscious, it seems particularly odd to permit a few instances of disunity to cut a swathe through the whole structure of consciousness, in the way that strong unity would require.[14]

[13] Consider, for example, Trevarthen's case, described in note 4.

*[14] See Michael Lockwood's very interesting discussion of the possibility of nontransitivity, in Lockwood (1989) chapter 6; on the present point, however, see especially p. 90. On the other hand, the suggestion that the structure of consciousness can be read off neuroanatomy is too simple. At least some distinctions with respect to neuroanatomical structure cut across

We need now to consider more carefully, however, whether the notion of weak unity can, even if we admit it, provide a way around the neo-Kantian argument from unity to objectivity. I have developed the notion of weak unity in pursuing what I called the naive objection to that argument. But I think we can now see that, interpreted this way, the naive objection fails. I'll first explain this in the abstract, then illus-

---

distinctions that it is natural to make with respect to the structure of consciousness. This can be seen by considering that when the corpus callosum is intact, we may have either a normal person with presumably normal unity of consciousness, or we may have a multiple personality patient who seems to support separate centres of consciousness with some duplication of contents. Moreover, when the corpus callosum is severed or absent, we may have either a commissurotomy patient, who seems to support separate centres of consciousness, or a callosal agenesis patient (born without a corpus callosum), who can pass almost all the experimental tests of unity that commissurotomy patients fail, even under conditions involving fixation. (On callosal agenesis see, for example, M. A. Jeeves 1965; A. D. Milner and M. A. Jeeves 1979; Stuart Diamond 1972, pp. 61–66.)

In this article I have been concerned to press a neo-Kantian argument to the effect that something objective must be appealed in order to resolve issues about the unity of consciousness. However, it is not easy to say just what that objective something might be without running into other difficulties. In *The Reappearing Self* I consider the way in which the attempt to read the structure of consciousness off neuroanatomy may involve the conflation of properties of vehicles of content with properties represented in content, criticized by Daniel Dennett, Ruth Millikan and others.

I there illustrate these considerations with a hypothetical case involving a congenitally acallosal person. Now in fact it is plausible to assume that acallosals depend on alternative internal mechanisms of informational integration to pass the tests that commissurotomy patients fail: ipsilateral and subcortical neural pathways. After all, they pass these tests under conditions of fixation, which should rule out the use of external mechanisms of integration such as cross-cuing between the hemispheres or access movements. However, we can consider a hypothetical acallosal who has always depended in ordinary circumstances on external rather than internal mechanisms of integration. I suggest that there is no reason in principle why such external mechanisms of information integration could not in ordinary circumstances support a unified consciousness. (See also Charles Marks 1981, at about p. 25, and note 42.) They might do so for an acallosal who has depended on them since infancy even though they would suggest communication between separate centres of consciousness in the case of a recently operated commissurotomy patient. I speculate that to place such a hypothetical acallosal under conditions of fixation, which prevent recourse to external mechanisms of integration, might be equivalent to a temporary surgical intervention: it might alter the structure of consciousness. Whereas to place a commissurotomy patient under conditions of fixation might merely block certain paths of information acquisition and communication between two separate centres of consciousness. Thus, intuitively, external mechanisms of integration might support a unified consciousness in some conditions but not others. Such a difference would need to be understood in objective terms, but evidently could not be read off neuroanatomy in any obvious way. I suggest that we should try to understand it in broadly functional terms, which draw on but are not limited to neurophysiological considerations. There is no reason in principle, that is, why the unity of consciousness could not be supported by causal mechanisms that pass outside of the central nervous system, so long as they do so in a way that meets the relevant functional criteria (which remain to be specified!).

trate the point by reference to the hypothetical based on the Sergent cases.

It is one thing to give up transitivity of co-consciousness of states at a time, so that the unity or separateness issue is not sharp and partial unity is possible. But it is another thing to give up the determinacy of such transitivity itself, so that what isn't determined is whether partial unity obtains or not. Indeed, the significance of giving up transitivity and admitting the third option of weak unity is largely undermined unless it can be determinately true that a case of weak unity obtains, a case showing co-consciousness not to be transitive. In trying to avoid the version of the neo-Kantian argument from subjectivity to objectivity that depends on agglomeration, we have painted ourselves into a corner. We admitted that determinacy of co-consciousness might be essential, but insisted it was directly open to consciousness and did not depend on indexed agglomeration or on transitivity. This gives rise to the possibility of merely weak or partial unity and nontransitivity. However, we can now say: even if we don't insist on transitivity of co-consciousness, surely it ought at least to be determinate whether or not a case of partial unity, a case demonstrating lack of transitivity, obtains. That is, given two thoughts, the thought that p and r and the thought that q and r, and given that there is no thought that p and q, it ought to be determinate whether there are two token conscious states of the same type, 'r' (duplication):

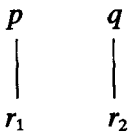

or rather there is just one conscious state of type 'r', which is co-conscious with a state of consciousness that p and also with a state of consciousness that q, even though the latter two states are not co-conscious (nontransitivity):

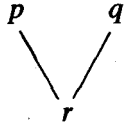

If the state of type 'r' is the same token in the two cases, and by direct appeal to consciousness we know that there is no thought that p and q, then we have a case showing the relation of co-consciousness to be nontransitive. But if a state of type 'r' is merely duplicated in the two

cases, then we do not have a such a case. Whether such nontransitivity obtains or not depends on whether a state of type '*r*' is duplicated, whether the '*r*'-type states are in fact distinct tokens. But that, surely, is not open to consciousness; if duplication obtains, it will not be apparent from the perspective of consciousness; the duplicates will not themselves be co-conscious. Whatever determines the identity or nonidentity of the token states of consciousness-that-*r* is not itself part of the content of consciousness. The types of conscious states may be open to consciousness, but the type-identity of conscious states occurring at the same time does not in general entail the identity of the relevant tokens of those types, as for example when you and I have conscious states of the same type at the same time. But then something objective is still needed, even if only to determine that nontransitivity as opposed to duplication obtains in a particular case.

Returning to the Sergent hypothetical, we see that the duplication option as well as the nontransitivity option is consistent with the type of evidence in question. It was natural to interpret the hypothetical by extending to the level of consciousness Sergent's suggestions about the limitation of disunity to certain types of information. But nothing so far rules out the other possibility, namely, the duplication of contents, so that each strongly unified half is separately conscious that the number on the right is higher than the number on the left, and there is no failure of transitivity. (See Figure 1.) What does the difference between the nontransitivity and duplication options amount to? There is no subjective viewpoint from which the issue can be determined. *If* it is determined, objective factors of some kind must determine it.

But another possibility is that in some cases there is simply no determinate answer to the duplication-or-nontransitivity question. That is, perhaps a kind of antirealism about the identities of token states of consciousness is warranted. Nothing I have said here argues against this possibility. But notice that this possibility goes beyond the one envisaged by Nagel in his early article on brain bisection and the unity of consciousness. There he considered whether in certain cases there might be no right answer to questions about how many centres of awareness there are. To admit this possibility is to admit the possibility of failures of transitivity, of partial or weak, as opposed to strong, unity of consciousness. But the further possibility we now consider, prompted by the question of what distinguishes cases of duplication from cases of nontransitivity, is whether in certain cases there may be no right answer to the question of whether weak or strong unity obtains.

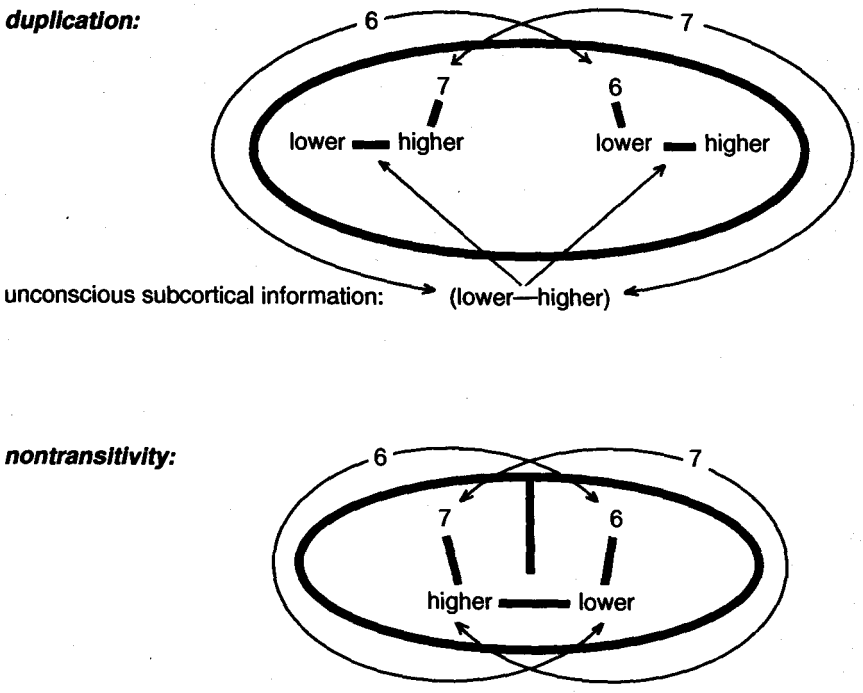

Figure 1. In the hypothetical example based on Sergent's cases, we need to distinguish *duplication* from *nontransitivity*.

7. Summary and concluding remarks

Let me summarize what I have argued so far, and then put it into context. The first neo-Kantian argument we considered insists on the need for some basis for the agglomeration of the contents of consciousness to reflect the strong unity of consciousness at a time. Since anything internal to content is 'just more content' for which the same need arises, the unity of consciousness requires an objective basis for agglomeration. Then I put a naive objection: why do we need a basis for agglomeration at all? Perhaps the relation of co-consciousness between conscious states at a time is not transitive, so that agglomeration fails and the strong unity of consciousness at a time does not obtain. But even so, a kind of weak unity of consciousness could still obtain, in that co-consciousness could still be determinate and open to introspection. In reply, I argued that while the co-consciousness of types of conscious state may be open to the perspective of consciousness, transitivity applies to token states. The issue between the dupli-

cation of tokens and nontransitivity in particular cases, such as the Sergent hypothetical, can't be determined from the perspective of consciousness, but requires objective determination as well. So, even the weak unity of consciousness at a time requires objectivity. Exactly *what* objective factors are needed to determine either strong or weak unity of consciousness, to the extent they are determined, is a further question, which has not been my focus here.

So, the neo-Kantian argument from the conditions of subjectivity to objectivity does not strictly need to assume indexed agglomeration or transitivity or strong unity of consciousness. The very identities of token states of consciousness themselves demand objectivity. We can short circuit the argument that the presence of consciousness demands the unity of consciousness, which in turn demands objectivity. What we have now got is an argument that making sense of the identity and distinctness of token states of consciousness requires something objective, something beyond the content of consciousness. In effect, we allow that the unity/separateness issue may not be sharply determined in the way the indexed agglomeration principle assumes. This undercuts the original just-more-content argument from unity to objectivity. But then, having admitted the possibility of weak unity, we demand at least that whether weak unity obtains or not should be determinate. And this in turn requires objectivity.

Now someone who is determined to avoid the need for objectivity, to defend a traditional view of consciousness as autonomously subjective, might be prepared to call the bluff of this neo-Kantian argument. That is, such a person may be prepared to admit that even whether weak unity obtains or not, hence the token identities of conscious states, may be indeterminate, in order to avoid the need for objectivity. I have already said that I have not argued against this move. In fact I don't want to argue against it, since it provides the other horn of a dilemma I have been aiming to set up for the traditional view of consciousness as autonomously subjective: it throws the baby out with the bathwater. From the traditional point of view, the price this move pays to avoid the need for objectivity, to keep the subjective realm autonomous, is very great: indeterminacy infects not just the unity of consciousness, but the very identity of conscious experiences and thoughts. The traditional view of the conscious realm as autonomous thus faces a dilemma: admit either this fundamental indeterminacy, or the need for objectivity. Either way, the traditional conception of consciousness has been partially deconstructed.

Someone might object to my construction of this dilemma as follows. There is no evidence for partial disunity in a Sergent-like case to begin with, just because there were not *simultaneous* manifestations of unity and disunity (compare section 4 of Marcel's commentary). But this objection puts too harsh a constraint on what might count as evidence of various facts about consciousness that hold simultaneously: facts about consciousness that hold simultaneously needn't necessarily be manifested simultaneously. In interpreting manifestations of facts about con sciousness, we are guided by certain background assumptions. One is that what it is possible for someone to be perceptually conscious of does not vary with his intentions alone, given that all sensory inputs are held constant over some relevant brief period of time; we can call this the *input assumption*. Another background assumption is that consciousness cannot be only partially unified; we can call this the *all-or-nothing assumption*. These two assumptions seem to conflict in a case like Sergent's, though the duplication option provides a way of avoiding the conflict.

In fact I don't think either of these assumptions is necessarily true. But what is wrong with the input assumption is not that it permits us to make inferences about what someone is aware of at one point in time from evidence gathered at a slightly earlier point. What is wrong with it is simply that it is false: what it is possible for someone to be aware of perceptually, and his perceptual experience, *may* vary with his intentions even when sensory inputs are held constant. The content of perceptual experience is not a function only of sensory input and 'upstream' processing, and in particular is not independent of the content of motor intentions. For example, when someone with paralysed eye muscles tries to glance to the right the world appears to jump to the right even though the pattern of light falling on the paralysed eye has not moved (see Gallistel, 1980, p. 175; Shebilske, 1984). More generally, experimental work in certain areas suggests that the nature of an intended response may alter the experience it is supposedly a response to.[15]

[15] See for example Richard Tegnèr and Maria Levander (1991); E. Bisiach, G. Geminiani, A. Berti and M. L. Rusconi (1990); E. Bisiach, A. Berti and G. Vallar (1985). Trevarthen's case, mentioned briefly in note 4, also raises issues about the input assumption, as indeed do the results reported by Marcel in his commentary. Whatever it is that I may be taking for granted about, as Marcel puts it, 'a single state of awareness of *one* aspect of *one* stimulus at a single time, in a single person' (p. 86), I do not accept the input assumption. Difference in conscious experience may depend on difference in response, whether the different responses in question are made simultaneously or successively.

It is the input assumption that puts interpretative pressure on the all-or-nothing assumption in the Sergent-like cases to begin with. And it was this pressure that led to the line of argument here developed and ultimately to the indeterminacy or objectivity dilemma for a traditional conception of consciousness as autonomous. If we give up the input assumption and allow that a change in intention between the same/different and higher/lower tasks in a Sergent-like case is sufficient for a change in perceptual consciousness despite constant sensory input, we may seem to have found a way of sidestepping the indeterminacy or objectivity dilemma I have tried to construct for the autonomy conception. However, what we have really done is to disturb an even more fundamental way of conceiving the mind, which generates the input assumption. This could be labelled the input-output picture, according to which the contents of perceptual experiences and intentions are independent:[16] perception is a matter of causal input, from the world to the mind, and intentional action is a matter of causal output, from the mind to the world. If the input-output picture of the relationship between the contents of conscious perception and intentional action is displaced, we can't just reinstate a traditional conception of the autonomy of conscious content. The displacement fundamentally restructures the ways in which issues about autonomy can be framed. But this is the beginning of another argument.

Note. I am grateful to the British Academy for its generous support of this work in the form of a Senior Research Readership (1990–1992). I am also grateful to many persons for comments on and criticisms of earlier drafts, but especially to John Campbell, Quassim Cassam, Derek Parfit, Nicholas Rawlins and Bernard Williams.

[16] Not instrumentally independent, of course: obviously, what one perceives has effects on what one decides to do, and one's intentional movements alter what one is in a position to perceive.

Bibliography

Allison, H. E. 1983: *Kant's Transcendental Idealism.* New Haven: Yale University Press.

Baron-Cohen, S., Leslie, A. M. and Frith, U. 1985: Does the autistic child have a 'theory of mind'? *Cognition*, 21: 37–46.

Bennett, J. 1966: *Kant's Analytic.* Cambridge: Cambridge University Press.

Bisiach, E., Berti, A. and Vallar, G. 1985: Analogical and logical disorders underlying unilateral neglect of space. In M. Posner and O. Marin (eds), *Attention and Performance*, vol. 11. Hillsdale, New Jersey: Erlbaum.

Bisiach, E., Geminiani, G., Berti, A. and Rusconi, M. L. 1990: Perceptual and premotor factors of unilateral neglect. *Neurology*, 40: 1278–1281.

Bisiach, E. and Luzzatti, C. 1978: Unilateral neglect of representational space. *Cortex* 14: 129–133.

Bisiach, E. and Vallar, G. 1988: Hemineglect in humans. In P. Boller and J. Grafman (eds), *Handbook of Neuropsychology*, vol. 1. Amsterdam: Elsevier.

Boden, M. 1990: *The Philosophy of Artificial Intelligence.* Oxford: Oxford University Press.

Brewer, B. 1992: Unilateral neglect and the objectivity of spatial representation. *Mind and Language*, 7: 222–239.

Cassam, Q. 1987: Transcendental arguments, transcendental synthesis, and transcendental idealism. *Philosophical Quarterly*, 37: 355–378.

Cassam, Q. 1989: Kant and reductionism. *Review of Metaphysics*, 43: 72–106.

Cassam, Q. forthcoming: Transcendental self-consciousness. In P. K. Sen and R. Verma (eds), *The Philosophy of P. F. Strawson.*

Chisholm, R. 1981: *The First Person.* Brighton: Harvester Press.

Collingwood, R. G. 1946: *The Idea of History.* Oxford: Oxford UniversityPress.

Coltheart, M. 1980: Deep dyslexia: a right-hemisphere hypothesis. In M. Coltheart, K. Patterson and J. C. Marshall (eds), *Deep Dyslexia*, pp. 326–380. London: Routledge and Kegan Paul.

Davidson, D. 1984: What metaphors mean. In *Inquiries into Truth and Interpretation*, 245–264. Oxford: Oxford University Press.

Davies, M. 1986: Tacit knowledge, and the structure of thought and language. In C. Travis (ed.), *Meaning and Interpretation*, 127–158. Oxford: Blackwell.

Davies, M. 1987: Tacit knowledge and semantic theory: Can a five per cent difference matter? *Mind*, 96: 441–462.

Davies, M. 1989: Tacit knowledge and subdoxastic states. In A. George (ed.), *Reflections on Chomsky*, 131–152. Oxford: Blackwell.

Dennett, D. 1984: Cognitive wheels: the frame problem of AI. In C. Hookway (ed.), *Minds, Machines and Evolution*, 129–151. Cambridge: Cambridge University Press.

Dennett, D. C. 1991: *Consciousness Explained.* Boston: Little, Brown.

Diamond, S. 1972: *The Double Brain.* London: Churchill Livingstone.

Evans, G. 1973: The causal theory of names. *Proceedings of the Aristotelian Society*, supp. vol. 47: 187–208.

Evans, G. 1981: Semantic theory and tacit knowledge. In S. Holtzman and C. Leich (eds), *Wittgenstein: To Follow a Rule*, 118–137. London: Routledge and Kegan Paul. (Reprinted 1985 in *Collected Papers*, 322–342. Oxford: Oxford University Press.)

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

Fogelin, R. 1985: *Hume's Skepticism in the* Treatise of Human Nature. London: Routledge & Kegan Paul.

Gallistel, C. R. 1980: *The Organization of Action: A New Synthesis.* Hillsdale, New Jersey: Erlbaum.

Gallistel, C. R. 1990: *The Organization of Learning.* Cambridge, Mass.: MIT Press.

Gazzaniga, M. 1988: In A. J. Marcel and E. Bisiach (eds), *Consciousness in Contemporary Science*, 226ff. Oxford: Oxford University Press.

Gibson, J. J. 1979: *The Ecological Approach to Visual Perception.* Boston: Houghton Mifflin.

Goldman, A. I. 1989: Interpretation psychologized. *Mind and Language*, 4: 161–185.

Goldman, A. I. 1992: In defense of the simulation theory. *Mind and Language*, 7: 104–119.

Goldman, A. I. 1993: The psychology of folk psychology. *Behavioral and Brain Sciences*, 16: 15–28.

Gopnik, A. and Wellman, H. 1992: Why the child's theory of mind really *is* a theory. *Mind and Language*, 7: 145–171.

Gordon, R. M. 1986: Folk psychology as simulation. *Mind and Language*, 1: 158–171.

Gordon, R. M. 1992a: The simulation theory: objections and misconceptions. *Mind and Language*, 7: 11–34.

Gordon, R. M. 1992b: Reply to Stich and Nichols. *Mind and Language* 7: 87–97.

Gordon, R. M. 1992c: Reply to Perner and Howes. *Mind and Language*, 7: 98–103.

Gordon, R. M. in press: Simulation without introspection or inference from me to you. In M. Davies and T. Stone (eds), *Mental Simulation: Philosophical and Psychological Essays.* Oxford: Blackwell.

Harris, P. L. 1989: *Children and Emotion: The Development of Psychological Understanding.* Oxford: Blackwell.

Harris, P. L. 1991a: The work of the imagination. In A. Whiten (ed.), *Natural Theories of Mind: The Evolution, Development and Simulation of Everyday Mindreading*, 283–304. Oxford: Blackwell.

Harris, P. L. 1991b: Letter to Josef Perner, 30 May 1991.

Harris, P. L. 1992: From simulation to folk psychology: the case for development. *Mind and Language*, 7: 120–144.

Heal, J. 1986: Replication and functionalism. In J. Butterfield (ed.), *Language, Mind and Logic*, 135–150. Cambridge: Cambridge University Press.

Heal, J. in press: How to think about thinking. In M. Davies and T. Stone (eds), *Mental Simulation: Philosophical and Psychological Essays*. Oxford: Blackwell.

Hurley, S. L. in preparation: *The Reappearing Self*.

Jeeves, M. A. 1965: Agenesis of the corpus callosum — physio-pathological and clinical aspects. *Proceedings of the Australian Association of Neurologists*, 3: 41–48.

Johnson-Laird, P. N. 1983: *Mental Models*. Cambridge: Cambridge University Press.

Kant, I. 1933: *The Critique of Pure Reason*. Tr. Kemp Smith, N. London: Macmillan.

Lockwood, M. 1989: *Mind, Brain and the Quantum: The Compound 'I'*. Oxford: Blackwell.

Marcel, A. J. 1993: Slippage in the unity of consciousness. In Ciba Foundation Symposium No. 174, *Experimental and Theoretical Studies of Consciousness*. Chichester: John Wiley.

Marks, C. E. 1981: *Commissurotomy, Consciousness and the Unity of Mind*. Cambridge, Mass.: MIT Press.

Milner, A. D. and Jeeves, M. A. 1979: A review of behavioural studies of agenesis of the corpus callosum. In I. S. Russell, M. W. Van Hof and G. Berlucchi (eds), *Structure and Function of Cerebral Commissures* 428–483. London: Macmillan.

Nagel, T. 1979: Brain bisection and the unity of consciousness. reprinted in T. Nagel, *Mortal Questions*. Cambridge: Cambridge University Press. (First published in 1971 in *Synthese*, 20.)

O'Keefe, J. 1985: Is consciousness the gateway to the hippocampal cognitive map? A speculative essay on the neural basis of mind. In D. A. Oakley (ed.), *Brain and Mind*, 59–98. London: Methuen.

O'Keefe, J. 1990: A computational theory of the hippocampal cognitive map. In J. Storm-Mathisen, J. Zimmer and O. P. Ottersen (eds), *Progress in Brain Research*, 83: 301–312. Amsterdam: Elsevier.

O'Keefe, J. 1991: The hippocampal cognitive map and navigational strategies. In J. Paillard (ed.), *Brain and Space*, 273–295. Oxford: Oxford University Press.

O'Keefe, J. 1993: Kant and the sea-horse. In N. Eilan, B. Brewer and R. McCarthy (eds), *Spatial Representation: Problems in Philosophy and Psychology*, 43–64. Oxford: Blackwell.

O'Keefe, J. and Nadel, L. 1978: *The Hippocampus as a Cognitive Map*. Oxford: Oxford University Press.

Parfit, D. 1984: *Reasons and Persons*. Oxford: Oxford University Press.

Peacocke, C. 1986: Explanation in computational psychology: language, perception and level 1.5. *Mind and Language*, 1: 101–123.

Peacocke, C. 1989: When is a grammar psychologically real? In A. George (ed.), *Reflections on Chomsky*, 111–130. Oxford: Blackwell.

Peacocke, C. 1992: *A Study of Concepts*. Cambridge, Mass.: MIT Press.

Peacocke, C. 1993: Externalist explanation. *Proceedings of the Aristotelian Society*, 93: 203–230.

Perner, J. 1991: *Understanding the Representational Mind.* Cambridge, Mass.: MIT Press.

Perner, J. and Howes, D. 1992: 'He thinks he knows': and more developmental evidence against the simulation (role taking) theory. *Mind and Language*, 7: 72–86.

Piaget, J. and Inhelder, B. 1951/1975: *The Origin of the Idea of Chance in Children.* New York: Norton.

Powell, C. T. 1990: *Kant's Theory of Self-Consciousness.* Oxford: Oxford University Press.

Quine, W. V. O. 1960: *Word and Object.* Cambridge, Mass.: MIT Press.

Rorty, R. 1970: Strawson's objectivity argument. *The Review of Metaphysics*, 24: 207–244.

Schwyzer, H. 1990: *The Unity of Understanding.* Oxford: Oxford University Press.

Sergent, J. 1990: Furtive incursions into bicameral minds. *Brain*, 113: 537–568.

Seymour, S., Reuter-Lorenz, P. and Gazzaniga, M. 1994: The disconnection syndrome: basic findings reaffirmed. Abstracted in The Society of Neuroscience, 1993.

Shebilske, W. L. 1984: Context effects and efferent factors in perception and cognition. In W. Prinz and A. F. Sanders (eds), *Cognition and Motor Processes.* Berlin: Springer-Verlag.

Shoemaker, S. 1984: Causality and properties. In S. Shoemaker, *Identity, Cause and Mind.* Cambridge: Cambridge University Press.

Sperry, R. W. 1990: Forebrain commissurotomy and conscious awareness. In C. Trevarthen (ed.), *Brain Circuits and Functions of the Mind.* Cambridge: Cambridge University Press.

Stich, S. and Nichols, S. 1992: Folk psychology: simulation or tacit theory? *Mind and Language*, 7: 35–71.

Stich, S. and Nichols, S. in press: Second thoughts on simulation. In M. Davies and T. Stone (eds), *Mental Simulation: Philosophical and Psychological Essays.* Oxford: Blackwell.

Strawson, P. F. 1959: *Individuals.* London: Methuen.

Strawson, P. F. 1966: *The Bounds of Sense.* London: Methuen.

Tegnèr, R. and Levander, M. 1991: Through a looking glass. *Brain*, 114: 1943–1951.

Trevarthen, C. 1974: Analysis of cerebral activities that generate and regulate consciousness in commissurotomy patients. In S. Dimond and J. G. Beaumont (eds), *Hemisphere Function in the Human Brain.* London: Elek Science.

Trevarthen, C. 1984: Biodynamic structures. In W. Prinz and A. F. Sanders (eds), *Cognition and Motor Processes.* Berlin: Springer-Verlag.

Walker, R. 1978: *Kant.* London: Routledge.

Wiggins, D. 1980: What would be a substantial theory of truth? In Z. van Straaten (ed.), *Philosophical Subjects: Essays Presented to P. F. Strawson*, 189–221. Oxford: Oxford University Press.

Wilkie, D. M. and Palfrey, R. 1987: A computer simulation model of rats' place navigation in the Morris water maze. *Behavioural Research Methods, Instruments and Computers*, 19: 400–403.

Williams, B. 1978: *Descartes: The Project of Pure Inquiry.* Harmondsworth: Penguin.

Wilson, M. D. 1987: *Descartes.* London: Routledge & Kegan Paul.

Wimmer, H. and Perner, J. 1983: Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13: 103–128.

Wimmer, H., Hogrefe, G.-J. and Perner, J. 1988: Children's understanding of informational access as a source of knowledge. *Child Development*, 59: 386–396.