

# 7

## Anosognosia, Schizophrenia, and Multiplicity

This chapter is concerned with a triad of clinical disorders in which the unity of consciousness appears to fall apart in fundamental ways. I begin with anosognosia, a syndrome in which patients are unaware of fundamental changes in the contents of their own conscious states. I then turn to schizophrenia, with a focus on two of its so-called ‘positive symptoms’—namely, thought disorder and thought insertion. The chapter concludes with an examination of dissociative identity disorder (or ‘multiplicity’), a condition in which individuals seem to manifest multiple conscious selves. As in previous chapters, my focus will be on the unity thesis: do these syndromes show that it is possible for human beings to have simultaneous states of consciousness that are not phenomenally unified? Along the way, we will pause to consider whether other forms of the unity of consciousness might be disrupted in these conditions.

### 7.1 Anosognosia: ‘blind to one’s blindness’

It was very striking that the patient did not take any notice of extreme and, later on, complete loss of vision. The patient, who was otherwise complaining a lot, was almost unaffected by this loss. When objects were presented in front of her, she—probably according to habit acquired during recent years—immediately tried to touch them, but she did not make an effort to recognize something by looking at it . . . It was obvious that she, like many blind people, had become an experienced guesser . . . She confirmed, calmly and faithfully, that she could see the presented objects, whereas almost daily examination proved the opposite (Anton 1899, trans. David et al. 1993: 267).

Anton called this impairment ‘soul blindness’ (*Seelenblindheit*), but it was a term coined by Babinski some fifteen years later—‘anosognosia’, literally, lack of knowledge of impairment—that stuck. (Curiously, Anton’s own name is reserved for a particular form of anosognosia, namely anosognosia for blindness.) Anosognosia can occur in the context of any number of deficits; it is possible to be anosognosic for paralysis, deafness, aphasia, and prosopagnosia

(the inability to recognize faces). Patients can be unaware of their physical impairments, of their behavioural consequences, and—what is most important from our point of view—even of the alterations to their own conscious states that result from those impairments and their behavioural consequences.<sup>1</sup>

These components of anosognosia are clearly exhibited by the following patient, an 80-year-old woman who was anosognosic for paralysis to the left side of her body due to a right-hemisphere stroke.

Examiner: Where are we?

C.C.: In the hospital.

Examiner: Which hospital?

C.C.: Santa Orsola.

Examiner: Why are you in the hospital?

C.C.: I fell down and bumped my right leg.

Examiner: What about your left arm and leg? Are they all right?

C.C.: Neither well nor bad.

Examiner: In which sense?

C.C.: They are aching a bit.

Examiner: Can you move your left arm?

C.C.: Yes, I can.

Examiner: [The examiner puts her right index finger in C.C.'s right visual field.] Can you touch my finger with your left hand? [C.C. does not move.]

Examiner: What happens?

C.C.: It happens that I am very good.

Examiner: Have you touched my finger?

C.C.: Yes.

Later in the interview:

Examiner: Could you clap your hands?

C.C.: I am not at the theatre.

Examiner: I know. But we just want to see whether you are able to clap your hands. [C.C. lifts her right arm and pits it in the position for clapping, perfectly aligned with the trunk midline, moving it as if it were clapping against the left hand. She seems perfectly satisfied with the performance.]

Examiner: Are you sure that you are clapping your hands? We did not hear any sound.

C.C.: I never make any noise. (Berti et al. 1998: 28–9)

<sup>1</sup> For reviews and cases see Bisiach & Berti (1995); Bisiach & Geminiani (1991); Forde & Wallesch (2003); Goldberg & Barr (1991); Heilman et al. (1998); Jehkonen et al. (2000); Marcel et al. (2004); McDaniel & McDaniel (1991); McGlynn & Schacter (1989); Papagno & Vallar (2003); Prigatano & Schacter (1991); Redlich & Dorsey (1945); Swartz & Brust (1984); Venneri & Shanks (2004); Vuilleumier (2000, 2004).

This transcript is representative of many involving patients with anosognosia. What might explain such reports?

Motivational accounts of anosognosia hold that patients are aware of their impairments and their consequences, but refuse to acknowledge them because their consequences are so overwhelming.<sup>2</sup> Although some cases of anosognosia may have a motivational component, it is implausible to suppose that motivational factors lie at the heart of the disorder. One problem with the motivational story concerns the selectivity of anosognosia. One patient may deny that they have hemianopia (blindness in one half of the visual field) but accept that they have hemiplegia (paralysis on one side of their body), whilst another might acknowledge their hemiplegia but deny their hemianopia (Nathanson et al. 1952; Jehkonen et al. 2000). Anton's own patient U.M. was 'blind to her blindness' but acknowledged her dysphasia. It would seem rather ad hoc to suppose that patients are less willing to acknowledge their blindness than they are their dysphasia.

Another account holds that although anosognosia does involve impairments of awareness, these impairments are limited to the encoding and retrieval of information in memory. On this view, patients are aware of their impairments when their behavioural consequences are obvious (say, because they fail to be able to dress themselves), but this awareness fails to make any impact on their long-term beliefs about their bodily state or capacities. This proposal receives some support from studies involving epileptic patients who have undergone the Wada procedure in preparation for temporal-lobe surgery. In this procedure, one hemisphere of the brain is anaesthetised, with the result that subjects suffer weakness in the contralateral side of the body. Because left-side injections produce language impairments in addition to right-side weakness, in many studies patients are questioned only after the effects of the barbiturate have resolved. A typical finding is that patients fail to remember that they had earlier experienced weakness in one half of their body. In one study, twenty-seven of thirty-one patients (who had had a right-sided injection) failed to recall having experienced left-arm weakness when questioned some fifteen minutes after the event (Carpenter et al. 1995; see also Meador et al. 2000).

Although this study suggests that memory problems might play a role in accounting for anosognosia, it is doubtful whether the condition can be fully explained in this way (Aimola Davies et al. 2008). Return to the transcript reproduced above. Not only did C.C. claim that she *could* move her arm and that she *had* moved her arm, she also claimed that she *was moving* her arm despite

<sup>2</sup> For classic statements of the motivational account of anosognosia see Weinstein & Kahn (1950, 1955). For recent defences of the approach see Ramachandran (1995) and Levy (2008).

the fact that it was immobile. As best one can tell, C.C.'s deficit wasn't merely one of failing to remember her impairment or its consequences, but included a problem of concurrent awareness.

Anosognosia appears to be a pathology of consciousness, but is it also a pathology of the *unity* of consciousness? Some authorities have thought so. Bisiach and Berti speculate that in anosognosia we might see a breakdown of phenomenal unity. 'Different mental states or events, although being individually endowed with phenomenal quality, are kept separate from one another within the stream of consciousness' (1995: 1338). The idea, I take it, is that patients are 'self-blind' to their impairments because certain components of their stream of consciousness have become dissociated from each other. The bulk of the patient's experiences remain introspectively accessible, but fragments of this stream have become isolated from the mechanisms responsible for introspection and cognitive influence more generally (see also Marcel 1993). This account might suggest that although C.C. retained some kind of awareness of her impairment (and perhaps also of its consequences), this state of awareness was not fully integrated with the rest of her conscious states. Let us call this the *phenomenal disunity* account of anosognosia.

Bisiach and Berti do not develop this proposal in much detail. As best I can tell, the only argument that they provide for it is an argument from 'dim' or 'clouded' knowledge—what Anton called 'dunkle Kenntnis'. Although anosognosia is characterized in terms of a lack of awareness of impairment, many patients possess an implicit appreciation of their impairment(s). For example, a patient with hemiplegia may deny that there is anything wrong with him, but he might nonetheless acknowledge that he doesn't retain the full suite of behavioural capacities that he once did. Other patients who also insist that there is nothing wrong with them will admit that the examiner would be unable to perform certain tasks had he or she been affected by the same impairment (Marcel et al. 2004). Still other hemiplegic patients might acknowledge their paralysis when directly questioned about it, but they will attempt to rise from bed or engage in other activities—such as knitting—that are obviously beyond them. Finally, when asked to account for their inability to carry out certain kinds of action, patients will often provide a justification that suggests that they have some awareness of their deficit. For example, a patient might explain her failure to move her arm by insisting that 'it has a cold' (Vuilleumier 2004).

I suspect that the argument from dim knowledge is meant to proceed along the following lines. The 'implicit awareness' of impairment that patients often manifest is best explained by supposing that there has been a division within the patient's stream of consciousness. The main tranche of this stream is available for the normal forms of cognitive and behavioural consumption (verbal report,

belief revision), but there is a subsidiary tranche—a tranche that contains a representation of the patient's impairment(s)—whose contents are able to exert only indirect influence on the patient's thought and action.

What should we make of this argument? The central problem with it concerns an assumption about how cognitive and behavioural control might be related to phenomenal fragmentation. The argument assumes that conscious states that are 'kept separate' from the rest of the subject's conscious states are likely to play only a restricted role in cognitive and behavioural control. It is not clear to me why we should accept this assumption. *Prima facie*, one would expect the contents of phenomenal states to be widely available for cognitive and behavioural control whether or not they are isolated from the bulk of the subject's conscious states. In short, it is not at all clear that phenomenal fragmentation provides any—let alone the *best*—explanation of dim knowledge.

So, how *should* we think of dim knowledge? That depends on the kind of dim knowledge in question. Some patients exhibit dim knowledge of their impairment in the sense that they are aware of certain aspects of it but not others. For example, a patient might seem to be oblivious to the impairment itself but aware of its behavioural consequences. This kind of dim knowledge does not seem to be particularly difficult to account for, at least in broad outline. The patients in question have simply failed to 'put two and two together'. They are aware of some facts about themselves (say, that they can no longer walk to the store), but not others (say, that one of their legs is weak). Admittedly, the facts of which they are unaware are intimately related to the facts of which they are aware—someone who is aware that they cannot walk to the store ought to wonder why this is the case—but anosognosic patients would not be alone in failing to draw out the implications of what they are aware of in the ways that they should.

However, other instances of dim knowledge are not easily explained by appealing to distinctions in the content of awareness. For example, a patient might appear to be aware of his impairment when required to produce a verbal report of it but not when required to engage in certain kinds of non-verbal behaviour. In such cases it looks as though the patient both is and is not aware of one the same fact. How might we explain this kind of dim knowledge?

The notion of probe-dependence might be of some assistance to us here (see §5.4) Consider a patient who is anosognosic for unilateral visual neglect due to right hemisphere damage. The patient might appear to be aware of her neglect when asked to copy a design but not when required to say whether anything is wrong with her. This dissociation may occur because verbal report activates the left hemisphere (which is responsible for speech production) and 'sucks' attention away from the damaged right hemisphere in which information about the

damage is represented. Without the demands of verbal report, attention can be restored to the damaged right hemisphere, and the patient will once again be conscious of the stimulus (Ricci & Chatterjee 2004). Attention can also be restored to the damaged hemisphere by caloric stimulation of the inner ear—either by stimulating the contralateral ear with cold water or the ipsilateral ear with warm water.<sup>3</sup> The probe-dependent perspective also suggests that even when a representation of the impairment does break through to consciousness it may do so only in a fairly weak and ineffectual manner, and as a result fail to possess the kind of influence on thought and action that is typical of conscious states.

Where does this leave the anosognosia-based challenge to the unity thesis? Not in particularly good shape, it seems to me. The central argument for thinking that patients with anosognosia might be subject to phenomenal fragmentation is the argument from dim knowledge, but we have seen that dim knowledge can be adequately accounted for in ways that involve no phenomenal disunity. Anosognosia is a pathology of consciousness, but it is not a pathology that puts any serious pressure on the unity thesis.

However, there is a sense in which patients with anosognosia might be said to have suffered from a loss in the unity of consciousness, broadly construed. The unity in question is not phenomenal unity but rather a kind of unity that is bound up with introspective awareness. Anosognosic patients are often oblivious to major changes in the contents of their own conscious states. They either overlook experiences that they have, or they fail to notice that they no longer have the kinds of experiences that they once did. For example, the anosognosic patient described above seems to be unaware of the fact that she was not clapping, despite the manifest difference in experience that typically accompanies the contrast between clapping and not clapping. Not only does this patient seem to be impaired in tracking what the world is like, she also appears to be impaired in tracking her own states of consciousness. Following Nikolidakos (2004), we might say that anosognosia involves a breakdown in the unity of *reflexive* consciousness.

In response to this proposal, a critic might challenge the assumption that the impairments seen in anosognosia will normally be introspectively manifest. Levine, for one, argues that the sensory impairments implicated in anosognosia are ‘never phenomenally immediate but instead must be discovered by observation and inference’ (1990: 234). This may be true of some of the impairments that occur in the context of anosognosia, but surely the fact that one is hemiplegic or blind ought to be apparent—indeed, manifestly apparent—to

<sup>3</sup> See Bisiach et al. (1991); Cappa et al. (1987); Rode et al. (1992); Rubens (1985); Silberpfennig (1949); Vallar et al. (1990); Vallar et al. (1994).

anyone with normal introspective abilities. How could anosognosia *not* involve some kind of impairment of the mechanisms underlying introspection?

One answer to this question takes the form of the ‘imagery hypothesis’, according to which patients are unaware of their impairments because the missing perceptual content has been replaced by internally generated experiences that ‘mimic’ the kinds of states that the patient expects and would have otherwise enjoyed (Heilman 1991; Goldenberg et al. 1995). Consider the following dialogue between a physician and H.S., a 46-year-old woman who appeared to be anosognosic for complete cortical blindness.

Examiner: (Moves a bunch of small keys, producing a sound): I am holding an object. Do you have any idea what it might be?

H.S.: Could that be a key?

Examiner: (Silently moves the keys beneath the table): What does it look like?

H.S.: On top there is a big ring, and it has a dark key-bit.

Examiner: Do you see the key well?

H.S.: I am seeing that it is a key.

Examiner: (Opens and shuts scissors): Do have any idea what that might be?

H.S.: Are those scissors?

Examiner: Do you see them?

H.S.: Only vaguely. I guessed a little.

Examiner: (Silently hides the scissors beneath the table): What can you see of these scissors?

H.S.: Upside are the handles where you take them, and below them is the part for cutting.

Examiner: Are you seeing this?

H.S.: Yes. (Goldenberg et al. 1995: 1378)

The authors of this paper suggest that H.S. had mistaken her imagery experiences for perceptual experiences, and that suggestion seems eminently reasonable. The objects that she described seeing were either the kinds of objects likely to be encountered in the environment in which she found herself, or—as in this extract—suggested by what she could hear. Arguably, H.S.’s blindness was not introspectively manifest to her because the missing perceptual content had been filled in by visual imagery, in something akin to the way in which the missing perceptual content in the blind spot is filled in by the mechanisms responsible for perceptual completion.

The imagery hypothesis has the resources to account for other forms of anosognosia as well. Consider again C.C., the hemiplegic patient described above, who—despite obvious evidence to the contrary—appeared to believe that she had just clapped. Perhaps C.C. had this false belief because she had the kind of proprioceptive and agitive experiences that normally accompany

clapping, although these experiences would not be 'perceptual' but 'imagistic'. Certainly some patients have such experiences. In response to a request to raise her arm a patient who had recently recovered from anosognosia for hemiplegia observed, 'It feels like it's rising, but it's not' (Chatterjee & Mennemeier 1996: 229). Such reports are not unusual in the literature.

But the imagery hypothesis does have its limitations. Consider the fact that C.C. *admits* that she's not making any sound when she 'claps'. Although she could be hallucinating motor experience, she does not appear to be hallucinating a sound, and not being able to hear yourself clap would alarm someone who was otherwise cognitively normal.<sup>4</sup> Similarly, one might have thought that the evident failure of H.S.'s visual experiences to support reliable action ought to have alerted her to the fact that they were internally generated rather than accurate representations of her environment (as she mistakenly believed them to be). So, even allowing for the fact that involuntarily generated imagery might plug the experiential gaps created by the patient's impairments, there is still a good case to be made for thinking that anosognosia may involve a 'breakdown' in the unity of reflexive consciousness.

Thus far I have considered cases in which patients appear to be unaware that they have *lost* certain forms of conscious experience. However, there are also patients who appear to be unaware that they *retain* certain kinds of conscious states. Consider the following case of 'inverse Anton's Syndrome' (Hartmann et al. 1991). The patient in question complained that he was completely blind, and took himself to be aware of the world only by non-visual means. One of his interests involved the buying and selling of horses. When asked how he judged the horses he said 'by their sound and their feel'. However, testing demonstrated that he had a relatively large area of preserved vision. He could name objects, colours, and famous faces that were presented in an area of 30 degrees in the right upper quadrant of his visual field. When confronted with his abilities he continued to deny that he had visual experience and insisted that he identified the stimuli in question by feel. The patient was neither disoriented nor hysterical, and the clinicians could not discern any plausible motive for his denial.

Was this patient conscious of various stimuli but merely unaware that he was conscious of them, or were his visual representations unconscious in the way in which the blindsight patient's visual representations appear to be? It is not clear. It is certainly tempting to contrast this case with that of blindsight, for unlike blindsighters this patient did not need to be prompted in order to use his visual

<sup>4</sup> I am indebted to Lizzie Schechter and Fiona Macpherson here.



information. Arguably, the most natural interpretation of the case is that the patient had conscious visual representations to which he lacked introspective access. He seems to have been aware of the features of the world that his visual experiences represented, but unaware of his visual experiences themselves. (To put it another way, he seems to have been unaware that his awareness of certain environmental features was visually mediated.) This self-blindness appears to have been selective, for there is no evidence that he was unaware of his experiences in other sensory modalities.

By way of drawing this section to a close, let me mention one implication that these findings might have for treatments of introspection. There are two broad perspectives one might take to introspection. On the one hand, one might think of introspection as a domain-general faculty (Schacter 1989; McGlynn & Schacter et al. 1988). On this view, selective failures of introspection would result from impairments to the connections between this faculty and particular perceptual systems. On the other hand, one might think of introspection as involving the operation of a number of domain-specific faculties. Furthermore, one might think of these domain-specific faculties as redeploying the very machinery used to generate experiences of the relevant domain in the first place (Bisiach 1988; Bisiach & Geminiani 1991). On this picture, the introspection of visual experience would involve a kind of reactivation of those experiences themselves. Following Bisiach, I submit that anosognosia provides some support for this domain-specific approach to introspection. Unlike the domain-general conception, the domain-specific model predicts that disruptions to a domain of experience will be accompanied by difficulties in the introspective detection of those disruptions. This is precisely what one finds in anosognosia: damage to certain modules not only impairs the individual's ability to form perceptual or sensory representations in a particular domain, it also impairs their ability to detect that impairment. This thought is actually hinted at by Anton himself, who remarked that in anosognosia for unilateral neglect 'not only the perception related to one side of the body or to one extremity but also the corresponding concepts are lost' (Anton 1898, trans. Förstl et al. 1993).

## 7.2 Schizophrenia: 'an orchestra without a conductor'

Another syndrome that is frequently said to involve a breakdown in the unity of consciousness is schizophrenia. The disease takes its name from the Greek for 'splitting of the psyche' and Kraepelin (1896), one of its earliest commentators, considered the loss of the unity of consciousness to be one of its core features.

Schizophrenia is a complex and heterogeneous disease, and many theorists have suggested that it is best thought of as a collection of loosely related disorders rather than a single entity.<sup>5</sup> Even those who argue in favour of retaining a single diagnostic category typically grant that the symptoms of schizophrenia cluster into two groups: negative symptoms and positive symptoms (Andreasen 1985; Crow 1980; Liddle 1987). Whereas the positive symptoms are characterized by a surfeit of affect and thought, the negative symptoms involve deficits in affect and thought, such as a reduction in the ability to experience pleasure (anhedonia), lack of motivation (avolition), and poverty of speech (alogia). Discussions of the structure of consciousness in schizophrenia typically restrict their attention to the positive symptoms, and it is on two such symptoms—namely thought disorder and thought insertion—that I will focus here.<sup>6</sup>

Thought disorder involves impairments in the ability to structure thought and action around goals. The cognitive and perceptual focus of patients is guided by associations and irrelevant stimuli that capture their attention, rather than by the logic of the task that they have set themselves. At root, thought disorder appears to involve a disturbance in selective attention.<sup>7</sup> Patients become ‘engrossed’ or ‘entranced’ by stimuli that should be disregarded. Attention is ‘sticky’, and cannot be easily moved from one stimulus to another:

If I am reading I may suddenly get bogged down at a word. It may be any word, even a simple word that I know well. When this happens I can’t get past it. It’s as if I am being hypnotized by it. It’s as if I am seeing the word for the first time and in a different way from anyone else. It’s not so much that I absorb it, it’s more like it is absorbing me. (McGhie & Chapman 1961: 109)

Another patient remarked that

the mind must have a filter which functions without our conscious thought, sorting stimuli and allowing only those which are relevant to the situation at hand to disturb consciousness. And this filter must be working at maximum efficiency at all times, particularly when we require a high degree of concentration. What happened to me . . . was a breakdown of the filter, and a hodge-podge of unrelated stimuli were distracting me from things which should have had my undivided attention. (MacDonald 1960: 218)

Whereas normal cognition is tuned to be sensitive to stimuli that are either intrinsically salient or important to the subject, in thought disorder attention is drawn to mundane matters—a stray remark, a vase, or the pattern on some

<sup>5</sup> See Andreasen & Carpenter (1993); Boyle (1990); Bentall (1990); and Poland (2007).

<sup>6</sup> See Freedman (1974) and McKay et al. (1996) for first-person perspectives on these symptoms.

<sup>7</sup> See Anscombe (1987) and Gray et al. (1991). However, see Cutting (1985) and McKenna (1994) for reservations about this account of thought disorder.

brickwork. The patient may be confused as to why her attention has been drawn to an insignificant stimulus, and in order to make sense of her own behaviour she might come to believe the stimulus *is* significant, that it possesses a hidden meaning that is manifest only to her.

As far as I can see, thought disorder does not put any real pressure on the unity thesis. There is no reason to think that patients have multiple conscious states that are not contained within an overall phenomenal field. What thought disorder does demonstrate is the importance of what we might call the *narrative* unity of consciousness—the ability to keep one’s thoughts on track. But narrative unity is not phenomenal unity, and the unity thesis is only concerned with the phenomenal structure of consciousness.

The breakdown in conscious unity exemplified by thought disorder is of a relatively recognisable kind, for we are all familiar with difficulties in screening out distracting information. A second positive symptom—thought insertion— involves forms of mental fragmentation that are rather more alien (Graham 2004; Mullins & Spence 2003). Jaspers characterizes thought insertion as follows:

Patients think something and yet feel that someone else has thought it and in some way forced it on them. The thought arises and with it a direct awareness that it is not the patient but some external agent that thinks it. The patient does not know why he has this thought nor does he intend to have it. He does not feel master of his own thoughts and in addition he feels in the power of some incomprehensible external force. (Jaspers 1962: 122–3)

Mellor gives the following oft-cited example of thought insertion:

I look out the window and I think that the garden looks nice and the grass looks cool, but the thoughts of Eamonn Andrews [a well-known TV presenter from the 1960s] come into my mind. There are no other thoughts there, only his . . . He treats my mind like a screen and flashes his thoughts into it like you flash a picture. (Mellor 1970: 17)

Before addressing the question of whether thought insertion might put any pressure on the unity thesis, we need to gain a clear conception of just what it involves. That turns out to be a far from straightforward task.

One account of thought insertion holds that it involves a loss of the sense of subjectivity or ‘my-ness’ that is said to accompany thought. On this view, patients are to be understood as denying that thoughts of which they are introspectively aware—thoughts that occur in their stream of consciousness—are theirs. This account is implicit in Freud’s description of thought insertion as a condition in which ‘portions of [the patient’s] mental life—his perceptions, thoughts, and feelings—appear alien to him and as not belonging to his own

ego' (Freud 1962: 13). Let us call this the *no-subjectivity* account. Another account of thought insertion holds that patients should not be understood as denying that they are the *subjects* of the thoughts in question, but merely that they are the *agents* of those thoughts (see e.g. Gallagher 2000; Stephens & Graham 2000). On this view, we should take 'the thoughts of Eamonn Andrews' to refer to the thoughts that Eamonn Andrews *produces*, not the thoughts that Eamonn Andrews *has*. As we shall see, this account can be developed in a number of ways.

The statements of the patients themselves do not clearly favour one model over the other. Although some expressions of thought insertion suggest that a sense of passivity lies at the root of the patient's experience—one patient may say that his mind is being treated like a screen; another may say that thoughts are put into his mind—other expressions of thought-insertion point in favour of the no-subjectivity model, insofar as these patients appear to deny ownership of the relevant thoughts and are inclined to attribute ownership of the thoughts to other individuals.

One might be inclined to reject the no-subjectivity account on the grounds that it ascribes to patients thoughts that are quite obviously fantastical—indeed, thoughts that may even be pragmatically self-defeating. But the force of this worry is uncertain, for patients with thought insertion are delusional, and by their very nature delusions involve gross departures from rationality. If it is possible for those with delusions of somatoparaphrenia to sincerely deny ownership of body parts that they admit are connected to their own bodies, why should it not also be possible for individuals with delusions of thought insertion to deny ownership over thoughts of which they are directly aware? As far as I can see, both the no-subjectivity and the no-agency accounts provide us with viable conceptions of what patients might mean in denying that certain thoughts are their own.

That being said, the no-agency account is *prima facie* more attractive than the no-subjectivity account, for it appears to make thought insertion *relatively* comprehensible. It is easier to understand why someone might come to think that their own thoughts are under the control of alien forces (and hence not their own in an agentive sense) than it is to understand how they might come to believe that their 'own' thoughts are not their own in a subjective sense. So, let us examine the prospects of the no-agency model.

Perhaps the most straightforward version of the no-agency account conceives of conscious thinking as an action that is typically accompanied by a 'feeling of doing'. According to this version of the account, it is the absence of this feeling that leads patients to suppose that their thoughts are not under their own control but are instead under the control of alien forces. Frith provides an influential presentation of this idea:

Thinking, like all our actions, is normally accompanied by a sense of effort and deliberate choice as we move from one thought to the next. If we found ourselves thinking without any awareness of the sense of effort that reflects central monitoring, we might well experience these thoughts as alien and, thus, being inserted into our minds. (Frith 1992: 81)

I think it is rather debatable whether conscious thought lives up to this description. In saying this I don't mean to deny that thinking *can* be accompanied by a sense of effort and deliberate choice; the experience of suddenly realizing that one's mind has been wandering and as a result setting oneself to concentrate on the task at hand is a familiar one. But, generally speaking, thought seems to be no more accompanied by a 'feeling of doing' than does perception or bodily sensation. For the most part one thought replaces another and the will seems to be a largely passive bystander. Even in the context of 'directed thinking'—which, it seems to me, is the exception rather than the rule—one does not experience oneself as the agent of particular thoughts but merely as having thoughts that are in line with one's overall cognitive goals.

Rather than develop the no-agency account in terms of 'a sense of effort and deliberate choice', we might do better to think of it in terms of a disruption to the intentional control of thought. We may not experience particular thoughts as 'things we do', but we do normally experience ourselves as having a certain kind of control over the general direction and evolution of our thoughts. We might call this kind of control 'metacognitive control.' Perhaps patients say that other agents are putting thoughts into their minds because they experience a lack of metacognitive control with respect to those thoughts.

The obvious objection to this proposal is that there are many contexts in which we take our thoughts to be inconsistent with our cognitive goals without ascribing them to other agents. Lying in bed at night I might curse my inability to control my thoughts, but I am not tempted to suppose that these thoughts are those of someone else. Moreover, as Stephens and Graham (2000) note, patients with obsessive thoughts experience their thoughts as escaping their metacognitive control, but they don't disown them in the way that patients with thought insertion do. An additional challenge for this account is that passivity phenomena are not restricted to judgements but include impulses and emotional states, so-called 'made emotions'.

I cry, tears roll down my cheeks and I look unhappy, but inside I have a cold anger because they are using me in this way, and it is not me who is unhappy, but they are projecting unhappiness into my brain . . . You have no idea how terrible it is to laugh and to look happy and to know that it is not you. (Mellor 1970: 17)

Unhappiness is not normally experienced as a state that is under one's control. Given this, it is hard to see how one might explain why a patient might come to believe that unhappiness is 'being projected into their brain' by supposing that they have experienced a disruption to the sense of metacognitive control.

A final variant of the no-agency approach can be found in Campbell (1999) and Stephens and Graham (2000), who suggest that thought insertion results from the patient's inability to reconcile the thoughts of which she is aware with her self-conception. As Stephens and Graham put it, the subject 'will not accept as agentically her own thoughts whose occurrence she finds inexplicable by reference to her conception or self-referential description of her intentional states' (Stephens & Graham 2000: 170).

In my view this is the most promising version of the no-agency account. Although it is not without problems—we often have thoughts that are neither explained by nor consistent with our self-conception without being at all tempted to alienate them—it does account for the fact that thought insertion can occur for mental states that we do not experience as under our direct control. However, this version of the no-agency view has departed so far from the origins of the approach that it is no longer clear that it is really an *alternative* to the no-subjectivity account. Suppose that one thought of oneself as an intentional system (Dennett 1971). On this conception, it is not clear that one could think of a mental state as one's own without also representing it as being caused by oneself in a certain kind of way. Arguably, belonging to an intentional system *just is* being produced by it. And if that is right, then perhaps this version of the no-agency account is best viewed not as an alternative to the no-subjectivity account but as a way of fleshing it out.

So much for the question of what might lie behind claims of thought insertion: do such claims—however they are to be understood—put pressure on the unity thesis? I think not. Although thought insertion reveals aspects of the unity of consciousness (broadly understood) that are easily overlooked, there is little reason to suppose that it involves any kind of phenomenal fragmentation. The aspect of the unity of consciousness that is lost or at least compromised in thought insertion is what we might call the unity of 'subjectivity' or 'reflexivity'—the sense that each of one's experiences are one's own. But—as we have remarked before (§1.3)—this form of the unity of consciousness ought not be confused with the phenomenal unity of consciousness.

The twin symptoms of thought disorder and thought insertion go some way towards justifying Kraepelin's description of 'the schizophrenic mind' (to use a dangerous phrase) as 'an orchestra without a conductor', but we have found nothing in our examination of schizophrenia that might call the *unity thesis* into

question. Certain forms of the unity of consciousness might be compromised in schizophrenia, but the kind of unity in which we are most interested here—phenomenal unity—appears to be left intact.

### 7.3 Multiplicity: ‘the delusion of separateness’

The clinical syndrome most closely associated in the popular mind with breakdowns in the unity of consciousness is multiple personality disorder (MPD), now officially known as dissociative identity disorder (DID). Both of these terms are cumbersome, and I will refer to the condition simply as ‘multiplicity’.

Multiplicity is sometimes confused with schizophrenia. The confusion is understandable, for there are certain parallels between multiplicity and the positive symptoms of schizophrenia (see p.169). Nevertheless, the disorders are regarded as distinct in clinical practice. Roughly speaking, schizophrenia involves a ‘fragmentation’ or ‘disintegration’ of the psyche, whereas multiplicity involves a ‘multiplication’ of the psyche (David et al. 1996). In order to qualify for a diagnosis of multiplicity, a person must have two or more distinct identities or personality states. These identities—or ‘alters’ as they are also known—‘each have their own relatively enduring pattern of perceiving, relating to and thinking about the environment and the self’ (American Psychiatric Association 2004). Alters take turns directing the behaviour of the multiple, and while a particular alter is ‘out’ the multiple’s behaviour will generally be guided only by the memories, beliefs, plans, and other intentional states of that alter.

The earliest reports of what might now be classified as cases of multiplicity were both published in 1791. The first involved a young man known to us only as ‘Captain Miller’s son’, who was described as having ‘two distinct minds, which acted by turns independently of each other’. When in dissociative ‘fits’ he could remember what had occurred during previous ‘fits’ but had no memory for the periods between fits (Carlson 1981). The second case involved a young German woman who suddenly exchanged her own personality for the manners and ways of a French-born lady. ‘In her French personality, the subject had complete memory of all that she had said and done during her previous French states. As a German, she knew nothing of her French personality’ (Ellenberger 1970: 127).

Interest in multiplicity flourished in the late nineteenth century, with theorists on both sides of the Atlantic—mostly notably Binet and Janet in France and James and Prince in the United States—arguing that dissociative phenomena are at odds with the claim that consciousness is necessarily unified (Crabtree 1986). This turn-of-the-century fascination with multiplicity reached its apex

with Morton Prince's (1905/1978) study of 'Christine Beauchamp' (whose real name was Clara Fowler) and W. F. Prince's (1915/16) study of Doris Fischer. After 1915, multiplicity enjoyed a dramatic decline; although some cases—such as Thigpen and Cleckley's (1954) 'Eve'—enjoyed a wide amount of attention, no more than fifty-four cases of multiplicity were reported between 1900 and 1970. By the 1970s 'double consciousness' had become 'multiple personality disorder', a change of name that was in part prompted by an increase in the average number of personality states manifested by 'multiples'; although multiple personalities were not unknown in the nineteenth century, the norm was two. Reports of duality are now infrequent, and the typical multiple is said to have between five to ten personality states, with some multiples reported to have hundreds of personalities.

How might multiplicity be at odds with the unity thesis? The most direct threat involves taking alters to have (or perhaps be) distinct streams of consciousness that might be 'out' at one and the same time. Let us call this the *phenomenal disunity* account of multiplicity. The phenomenal disunity account is at odds with the unity thesis, for if the unity thesis is right then the conscious states enjoyed by a subject at any one time must be phenomenally unified with each other. Of course, one could attempt to reconcile the phenomenal disunity account of multiplicity with the unity thesis by identifying subjects of experience not with human beings but with (say) psychological networks or intentional systems. (Indeed, on this view the advocate of the unity thesis might even *expect* alters to 'have' their own stream of consciousness.) Just what to say about the connection between the unity of consciousness and the self should we allow that a single organism might 'house' multiple subjects of experience raises many complex issues, not least of which is the problem of how to individuate intentional systems. However, I will leave these issues to oneside here and focus solely on the question of whether the phenomenal disunity account of multiplicity is correct. If it is not—as, indeed, I shall argue is the case—then these issues are moot.

Although discussions of multiplicity often assume that the phenomenal disunity account is true, the view has rarely been defended. In fact, the only developed defence of the view that I know of is to be found in Stephen Braude's book *First-Person Plural*. Braude's account of multiplicity is nuanced. Although he argues that alters have their own streams of consciousness—indeed, that they are independent loci of *self*-consciousness (1995: 78 f.)—he also argues that multiples have a single, underlying self, which he describes as a 'Kantian ego'. I will leave the Kantian components of Braude's account to oneside, and focus only on his arguments for the claim that alters possess autonomous streams of consciousness, streams that can run in parallel to each other.



On my reading *First-Person Plural* contains four arguments for the claim that alters are distinct loci of consciousness. The first of these arguments appeals to the intentional disunity exhibited by alters. Not only do alters claim to be

distinct persons (and not merely personalities); they will disavow the interests and activities of other alters and pursue goals of their own, sometimes even holding different jobs. Moreover, different alters seem to be of different ages and sexes and appear to have distinct overall body-images. For example, one alter might feel he or she is the wrong sex, or too young, short, or fat, to wear the clothes of another alter. (1995: 67; see also 253)

Call this ‘the argument from intentional disunity’. The argument starts from a firm foundation, for alters do indeed possess their ‘own’ traits, beliefs, memories, goals, and self-conceptions. The problem arises when we attempt to move from that claim to the claim that alters are distinct loci of consciousness. Should we think of these sets of psychological states as organized around ‘discrete centres of self-consciousness’?

I think not. Although some authors have suggested that ‘the grounds for assigning several selves to [a multiple] can be as good as—indeed the same as—those for assigning a single self to a normal human being’ (Dennett & Humphrey 1998: 54, emphasis suppressed), in my view the urge to reify alters in this way should be resisted. Alters ought to be regarded as personality ‘states’ or ‘files’ rather than bona fide subjects of experience. It might be tempting to argue that since by definition any multiple will have a number of self-conceptions, he or she will also have a number of conscious selves. However, we can see that this inference is fallacious by considering self-deception. Jane has good reason to believe that her husband is having an affair, yet she ‘keeps the truth from herself’, as we say. When Jane is being honest with herself she recognizes that her marriage has fallen apart, but when in the grip of self-deception she thinks of herself as someone who is happily married. Jane has distinct self-conceptions, but there is no reason to regard these self-conceptions as picking out different selves in any robust sense of the term. Similarly, we can regard alters as self-conscious without being committed to the view that the first-person thoughts ‘had by’ a multiple’s various alters refer to distinct entities.

Self-deception not only undermines the argument from intentional disunity, it also provides us with a lens through which to view multiplicity: perhaps multiples are simply massively self-deceived. By this I do not mean that multiples are deceived *by* themselves (although this may indeed be true), but rather that they are deceived *about* themselves (Heil 1994). In one alter state a white middle-aged male multiple might believe that he is black; in another alter state he might believe that he is a woman; and in a third alter state he might take himself to be a child. As Putnam et al. (1986) put it, multiplicity is a ‘delusion of

separateness'. The most striking manifestation of this delusion is 'internal homicide', in which one alter tries to kill a fellow alter, oblivious to the fact that any such act would bring about his or her own demise.

Even if it were sound, the argument from intentional disunity could at best establish only that multiples have *successive* rather than *simultaneous* streams of consciousness—in the terminology of the nineteenth century, that patients have *alternating* rather than *double* consciousness. Braude's second argument, if successful, would make good on this lacunae. The argument appeals to certain features of the way in which multiples 'switch' between alter states:

one can actually observe and clearly identify the participants in the struggle. For example, as two alters vie for executive control, the multiple's face might shift rapidly between the distinctive features of each. Even more importantly, the clear personality shifts on the subject's face often reflect the alters' idiosyncratic contributions to the conflict. For example, one personality might show anger, tension or confusion, and the other might display amusement and contempt. And those dispositions can be exhibited in a manner characteristic of the respective personalities. (Braude 1995: 67 f.)

The idea, I take it, is that multiples harbour multiple *agents*, and where we have multiple loci of agency it is likely that we also have multiple loci of consciousness.

Even if switching lives up to Braude's description of it—see Hacking (1995) for a rather different picture—I suspect that the conflict Braude describes is merely an exaggerated form of the struggle for emotional control with which many of us are familiar. Consider a person who has been deeply insulted in a context in which anger is not an appropriate emotion to manifest. One might witness a struggle between anger and self-control being played out on the subject's visage. Braude's comments suggest that inter-alter conflict is 'deeper' than this, but I am not convinced. For one thing, the agentic disunity described by Braude and other commentators appears to be superimposed on a base of sensorimotor integration. To the best of my knowledge, multiples do not demonstrate the kind of agentic disunity seen in (say) the anarchic hand syndrome.

Even if there is a sense in which alters vie for control of the multiple, it doesn't follow that we should conceptualize switching in terms of a struggle between two centres of apperception—two self-conscious agents. Rather than describing *alters* as vying for executive control, it might be more perspicuous to describe the multiple's behaviour as successively informed by competing intentional structures. Indeed, in some passages Braude describes *alters* as switching (as though alters themselves are loci of agency), while in other passages he describes *the multiple* as switching between alter states. It is this latter locution

that seems to be most common in the experimental literature. In a representative study, patients were ‘required to select two personalities unaware of each other’ (Dorahy 2001: 777; see also Loewenstein et al. 1987). It is not clear to me how the patients could comply with this request if it is their alters—rather than ‘they themselves’—that are the basic units of agency.

A third argument for phenomenal disunity—perhaps more suggested by what Braude says than endorsed as such—concerns the role that switching might play within the life of the multiple.

Switching personalities enables a multiple to cope with exhaustion, pain, or other impairments to normal or optimal functioning. For example, if A is tired or drugged, B can emerge fresh or clear-headed. When in pain, A can switch to an anesthetic personality. Or, personalities can keep passing the pain to each other in turn, switching when the persistent pain becomes intolerable. (Braude 1995: 45)

Braude’s comments have their roots in a common account of the aetiology of multiplicity, according to which multiplicity arises out of the attempt to cope with the psychic pain occasioned by horrific abuse. The victim of the abuse deals with the pain by creating other personalities to whom it can be transferred. The thought behind this passage, I take it, is that switching could play this role only if alters qualify as distinct loci of consciousness.

This proposal might contain a kernel of truth, but it is difficult to make sense of if taken literally. Pains are not the sorts of things that can be passed from one subject of experience to another. I can *cause* you to be in pain, but I cannot *give* you my pain in the way in which I can give you my sandwiches, my shoes, or even the shirt off my back. Moreover, even if pains were transferable, we have no conception of *how* they might be passed from one alter to another. Nor is it clear how the multiple might be better off by transferring pain between their alters. Wouldn’t the multiple him- or herself still be in pain irrespective of which of their alters ‘had’ the pain?

However we can salvage something from the proposal. Drawing on an approach that dates back to the work of Theodule Ribot (1891), suppose that we think of alters as behavioural schemas—networks of intentional states that govern an organism’s responses in particular environments.<sup>8</sup> Behavioural schemas are not unique to multiples, but structure all our interactions with the environment. Switches between one schema and another can be triggered by changes in environment, as when a teacher takes on a pedagogical persona upon entering a classroom, but they can also be endogenously elicited, as when one adopts a certain mood state in order to cope more effectively with a

<sup>8</sup> See also Putnam (1986); Bower (1994); Silberman et al. (1985).

challenging situation. What marks out multiplicity as a form of pathology is the fact that the patient's schemas are abnormally insulated from each other, and also the fact that the multiple's schemas frequently contain delusional content. The multiple often deals with her environment by taking on a schema that misrepresents her true identity.

On this picture of things, what it is for one alter to 'transfer' its pain to another is just for the multiple to switch from one schema state to another. And now we can understand how the 'transfer' of pains might be of benefit to the patient, for some personality states might be better equipped to deal with noxious stimuli than others. We know that certain types of pain can be ameliorated by various cognitive strategies, and the patient might be better at implementing such strategies when in some personality states than in others. Indeed, by switching from one alter state to another the multiple might not only be able to handle psychic distress more effectively but in fact avoid it altogether.

We have made sense of the idea that alters can 'transfer' their pains, but in so doing we have also deflated any hope that the argument might have established phenomenal disunity. The transfer of a pain from one alter to another involves a single stimulus being processed within the context of distinct behavioural schemas, rather than the movement of a single conscious state from one stream of consciousness to another. There is no evidence of phenomenal disunity here.

Braude's fourth and final argument for the view that multiples have multiple streams of consciousness appeals to introspection. Consider the following quotation, taken from W. F. Prince's description of his patient Doris Fischer. The narrator is one of Fischer's alters, Sleeping Margaret [S.M.], and the initials refer to her other alters.

S.D. [Sick Doris] watched when R.D. [Real Doris] was out. There would be three of us watching her, each with thoughts of her own. S.D. watched R.D.'s mind, M. [Margaret] watched S.D.'s thoughts of R.D., and I watched all three. Sometimes we had a disagreement. Sometimes a jealous thought would flit through S.D.'s mind—she would think for a moment that if R.D. would not come out any more M. might not like her (S.D.) as well as R.D. She never tried to hinder R.D.'s coming out, though, but always to help, and only a slight thought of the kind would flit through her mind. But M. would see it and get cross with S.D., and so the disturbance inside would make R.D. go in. (Prince 1915/16: 109)

The phenomenon referred to here in which one alter appears to be directly aware of the thoughts of a co-alter is often known as 'co-consciousness'. However, this term is less than ideal, for not only is 'co-consciousness' often used as a synonym for 'phenomenal unity', the term implies that the kind of access in question is symmetrical, which is not the case. Indeed, this kind of

access is typically asymmetrical, with alter A having access to alter B's mental states but not vice versa. For lack of a better term, I will call this relation 'inter-alter access'.

Inter-alter access is often presented as a kind of telepathy, as if alters who enjoy it are 'able to peek into a private room of experiences, or access or "read" a stream of experiences distinct from their own' (Braude 1995: 82).<sup>9</sup> The idea, I take it, is that alters occasionally have introspective (or, if you like, 'quasi-introspective') access to two kinds of mental states: their own and those of certain co-alter. In the Doris Fischer case, the narrating alter (S.M.) would have introspective access both to her own thoughts and to those of S.D., R.D., and so on. Not only would S.M. be aware of *what* each of these alters is thinking, she would also be aware of *which* particular alter was thinking each of the various thoughts to which she had 'quasi-introspective' access. For example, S.M. would know that a certain thought was flitting through S.D.'s mind rather than through (say) R.D.'s mind.

Although the telepathic model of inter-alter access has a certain charm, there is no shortage of objections to it. How might introspection go about tagging thoughts as the thoughts of particular alters? Why might such a mechanism have evolved? How could one be introspectively aware of a thought without being aware of it as one's own? It is far from clear that there are good answers to any of these questions. We might be forced to endorse the telepathic model even in the face of these challenges if it were the only game in town, but it isn't. In fact, there are a number of other ways in which inter-alter access can be conceptualized.

One alternative to the telepathic account holds that reports of inter-alter access are confabulations—'mere hallucinations'—of mental states (Stephens & Graham 2000). The introspective state that the multiple is reporting might be real enough, but the mental state that is its target might be a figment of the multiple's imagination. Support for this proposal is provided by the fact that alters can be created in hypnotic contexts as merely intentional entities (Merskey 1992). Pierre Janet describes the evolution of a secondary personality in a woman, called Lucie, whom Janet was treating for fits of terror.

'Do you hear me?' asked Janet.  
 'No,' she answered (in writing).  
 'But you have to hear in order to reply.'  
 'Yes, of course.'

<sup>9</sup> For other 'telepathic' treatments of inter-alter access see Rovane (1998); Greenwood (1993); Wilkes (1988); and Zemach (1986).

‘Then how do you do it?’  
 ‘I don’t know.’  
 ‘Must there not be someone who hears me?’  
 ‘Yes.’  
 ‘Who is it?’  
 ‘Someone other than Lucie.’  
 ‘Oh, indeed. Another person. Should we give this person a name?’  
 ‘No.’  
 ‘Yes. It is more convenient.’  
 ‘All right then—Adrienne.’  
 ‘Adrienne, do you hear me?’  
 ‘Yes.’ (Janet 1913: 318)

Braude (1995: 25) describes Janet as using automatic writing to discover Adrienne, but it seems more fitting to describe Janet as ‘aiding materially in the formation of a person’, as Binet (1890/1977b) put it in his rather understated way. Here, as elsewhere, fiction may give rise to fact. Alters might begin life as purely intentional entities, figments of hypnotic hallucination, but thereafter acquire a degree of reality as the multiple begins to live out her fantasy (Velleman 2006).<sup>10</sup>

Another alternative to the ‘telepathic’ account holds that in inter-alter access multiples are aware of genuine mental states, but these states are their own rather than those of some other subject of experience. This proposal receives some support from the fact that the line between inter-alter access on the one hand and the schizophrenic symptoms of thought insertion and auditory hallucination on the other is far from sharp (David et al. 1996). Patients with schizophrenia are often aware of thoughts or voices that they take to be self-referential (‘She’s so mean...’), and it is not hard to see how such states might be taken as manifestations of inter-alter access by a sympathetic theorist. Bliss et al. (1983) studied forty-four patients with auditory hallucinations, thirty-five of whom had received a diagnosis of schizophrenia at some point. Of these thirty-five patients, the authors identified twenty as having multiple personality disorder on the grounds that they manifested multiple personalities under hypnosis—‘the voices could be contacted, engaged in conversation, and would readily admit to being the culprit’ (1983: 30). I very much doubt that this study shows what its authors take it to show—namely, that many individuals with multiplicity are falsely diagnosed with schizophrenia—but it does suggest that the distinction between auditory hallucinations and inter-alter

<sup>10</sup> See Braude (1995); Harriman (1942); Kampman (1976); and Putnam (1986) for discussion of the creation of alternate personalities in hypnotic contexts.

access is neither sharp nor easily discerned. This point is reinforced by other studies of dissociative identity disorder. One study of thirty patients found that thirty per cent of the patients heard voices commenting on their actions and forty-three per cent took the thoughts of others to be inserted into their minds (Kluft 1987); another study of 102 patients found that seventy-nine per cent of patients reported voices commenting on their actions while sixty-five per cent reported experiencing thoughts that they ascribed to others (Ross et al. 1990). Bizarrely, this latter study also found that those patients who had received a diagnoses of dissociative identity disorder had *higher* levels of first-rank symptoms of schizophrenia than had those patients who had been diagnosed as suffering from schizophrenia, a finding that has been confirmed by other studies (Steinberg et al. 1994).<sup>11</sup>

What does this show? Well, we have little hesitation in regarding the inner speech and thoughts of which the subject is apparently aware in auditory hallucinations and thought insertion as actually belonging to the patient, notwithstanding her protestations to the contrary. The patient is talking or thinking to himself, unaware that this is what she is doing. Why shouldn't we say precisely the same thing about inter-alter access? Why should we not say that S.M. is aware of her own thoughts and is simply mistaken in ascribing them to another 'alter'?

The obvious response is that there is a sense in which multiples *do* have multiple minds whereas patients with schizophrenia do not. Although there is something to this response, I don't think that it undermines the force of this proposal. Returning to the Doris Fischer case, we can ask whether S.M.'s reports of being aware of S.D.'s jealous thoughts 'answered' to anything. They might, if Prince's account of the case is to be believed. Doris Fischer had an alter called (by Prince and herself) 'Sick Doris', and she may well have had jealous thoughts while in this alter state. But this does not show that S.M. had any kind of 'telepathic access' to S.D.'s mind, for there are many 'third-person' ways in which Doris Fischer could, as S.M., divine that she had an alter called S.D. with jealous thoughts. Most obviously, she could have acquired this information from her therapist, Prince. Having represented herself as having an alter with jealous thoughts, Doris Fischer now has reason to create such an alter if she wants to be seen by Prince as a good patient. Moreover, it is entirely possible that she doesn't want to acknowledge or identify with those jealous thoughts that she might have when S.M. is 'out', and hence ascribes them to another alter—'*I'm* not jealous, S.D. is'.<sup>12</sup>

<sup>11</sup> As Kennett and Matthews (2003) note, there are also deep commonalities between inter-alter access and the phenomenon of depersonalization, a condition in which patients experience themselves as alienated from their own mental states (see Chapter 11).

<sup>12</sup> Thanks to Ian Phillips here.

A final account of inter-alter access—in some ways a variant of the model just discussed—is that inter-alter access involves the mental states of one alter ‘leaking’ into those of another alter. Let me explain, I have suggested that we should conceive of alters as ‘psychological schemas’: semi-autonomous clusters of behavioural traits, dispositions, beliefs, memories, and other intentional states. Suppose that while Doris Fischer is in one alter state a thought that ‘belongs to’ another alter is activated for some reason. The intruding thought will seem alien to her, for it won’t sit well with her self-image—that is, the self-image that is currently structuring her overall thought and action. And if indeed the thought doesn’t ‘belong’ to the psychological schema that is currently directing her behaviour then there may be a sense in which she is right to reject it as not fully hers.

Something akin to this phenomenon may not be uncommon in everyday life. Suppose that one is in a sexually intimate context when thoughts appropriate to a very different context—say, a philosophical discussion—suddenly enter one’s head. Such thoughts might seem alien, as not truly one’s own. If one also had a model of oneself as having a philosophical personality for which such thoughts would be appropriate, one might be tempted to ascribe such thoughts to that personality and not one’s current personality. In this respect it is interesting to note that the kind of states for which inter-alter access is reported tend to be affectively laden. Perhaps this is because it is affectively laden states that are most at odds with an alter’s self-conception—that is, with ‘its’ conception of the multiple’s identity.

I conclude that the argument from inter-alter access fails to support disunity interpretations of multiplicity. Where does this leave Braude’s defence of the phenomenal disunity treatment of multiplicity? Braude’s case in favour of the disunity account of multiplicity is advanced as a cumulative one, and it is entirely possible that although none of his arguments is individually convincing his overall case is. I will leave readers to judge for themselves whether that might be so, but my own view is that the unity thesis is unscathed by what we know of multiplicity.

## 7.4 Conclusion

The disorders of consciousness seen in anosognosia, schizophrenia, and multiplicity are frequently said to involve a breakdown in or fragmentation of the unity of consciousness. There is *some* truth in such claims given that certain elements of the unity of consciousness—broadly construed—are impaired in each of these syndromes. In anosognosia we see a breakdown in the integration that



normally holds between the contents of a person's consciousness and their introspective access to those contents. In schizophrenia we see a disruption to the narrative unity of consciousness brought about by impairments to selective attention in the context of thought disorder, and a disruption to the unity of self-consciousness brought about by the patient's inability to keep track of their own in the context of thought-insertion. Multiplicity presents us with an even more profound impairment to the unity of self-consciousness. Here, the patient suffers from the delusion that he or she is in some way multiple subjects of experience. In varying ways, each of these syndromes involves notable departures from the coherence and integration that consciousness—particularly *self*-consciousness—normally displays.

But none of these syndromes threatens the unity thesis. In none of these cases do we have good reason to posit a breakdown in the *phenomenal* unity of consciousness. The kinds of failures in representational integration and co-accessibility of content that are required to construct a plausible case against the unity thesis are notably missing from the clinical literature on these syndromes. Perhaps the strongest evidence against the unity thesis found thus far derives from the behaviour of individuals with multiplicity. But even here, I have suggested, it is more plausible to suppose that the agentic disunity seen in multiplicity is best accounted for in terms of a single stream of consciousness that is successively informed by a variety of psychological schemas, rather than by appeal to parallel streams of experience. This switching between schemas may produce the *appearance* of phenomenal disunity, but I suggest that this appearance masks an underlying phenomenal unity.