

Shaun Gallagher and Anthony J. Marcel

The Self in Contextualized Action

Abstract: *This paper suggests that certain traditional ways of analysing the self start off in situations that are abstract or detached from normal experience, and that the conclusions reached in such approaches are, as a result, inexact or mistaken. The paper raises the question of whether there are more contextualized forms of self-consciousness than those usually appealed to in philosophical or psychological analyses, and whether they can be the basis for a more adequate theoretical approach to the self.*

First, we develop a distinction between abstract and contextualized actions and intentions by drawing on evidence from studies of rehabilitation after brain damage, and we introduce the notion of intentional attitude. Second, we discuss several interesting conclusions drawn from theoretically and experimentally abstract approaches. These conclusions raise some important issues about both the nature of the self and reflexive consciousness. At the same time they indicate the serious limitations concerning what we can claim about self and self-consciousness within such abstract frameworks.

Such limitations motivate the question of whether it is possible to capture a sense of self that is more embedded in contextualized actions. Specifically, our concern is to focus on first-person approaches. We identify two forms of self-consciousness, ecological self-awareness and embedded reflection, that (1) function within the kinds of contextualized activity we have indicated, and (2) can be the basis for a theoretical account of the self. Both forms of self-consciousness are closely tied to action and promise to provide a less abstract basis for developing a theoretical approach to the self.

To get clear about philosophical problems, it is useful to become conscious of the apparently unimportant details of the particular situation in which we are inclined to make a certain metaphysical assertion.

Wittgenstein (1958)

The self that we are does not possess itself; one could say that it 'happens'.

Gadamer (1976)

Overt action is indivisible . . . it is the whole individual who acts in the real environment.

Neisser (1988)

Correspondence [Co-equal authorship; names in alphabetical order]: Shaun Gallagher, Philosophy and Cognitive Science, Canisius College, Buffalo NY 14208, USA. Email: gallaghr@canisius.edu
Anthony J. Marcel, Medical Research Council, Cognition and Brain Sciences Unit, 15 Chaucer Road, Cambridge CB2 2EF, UK. Email: tony.marcel@mrc-cbu.cam.ac.uk

Surprising and seemingly counter-intuitive results are not uncommon when philosophers, psychologists, and neuroscientists, employing a variety of first- and third-person approaches, search for an adequate model of the self. At least one philosopher equates the self with a momentary existence so that we are said to live through a large number of consecutive momentary selves (Strawson, 1997). Other philosophers, introspectively exploring the stream of consciousness, fail to find anything at all that resembles a self (Hume, 1739). Psychological and neurological observations and experiments (concerning, for example, Dissociative Identity Disorder and split-brain phenomena) suggest the possibility of more than one self to a single human organism. (Radden, 1998; Gazzaniga, 1978; Sperry, 1968a,b). Neuroscientists seem quite intent on demonstrating that what we call the self is either nothing more than a set of neuronal processes (Crick, 1994) or what such a set of processes produces (Ramachandran and Hirstein, 1997). Others think that the various ideas of the self as unitary, unique, familiar, autonomous, and so forth, attributable to common sense or to Enlightenment or Romantic conceptions, amount to naive delusions, convenient defence mechanisms, or, at best, abstract centres of narrative gravity.¹

When faced with a range of questions about self (questions pertaining to identity, experience of self, nature of self, and so forth) most theorists approach the topic in a manner that is abstract or detached from behaviour and action normally embedded in pragmatically and socially contextualized situations. When, for example, philosophers employ reflective introspection in order to search for the unity of consciousness or ‘the self’ as an element in consciousness, they choose a framework for their investigation that is not equivalent to the framework within which people normally act. The introspective framework takes consciousness and the self as objects and thereby fails to capture their role in the realm of action, where they are specifically not objects. Similarly, psychological experimentation sometimes places subjects in circumstances where they are called upon to view their own body or their own thinking processes in an abstract and detached way. We argue, in this paper, that these various approaches to developing a model of the self, either through methods of reflective self-consciousness or by means of scientific experimental investigation, have been conducted from perspectives that remain relatively abstract in a way that disqualifies, or at the very least places qualifications on many of these findings. We want to define what these perspectives have in common, that is, in what sense they are abstract. We also want to suggest that most of the controversies, problems and paradoxes concerning the notion of self are the result of searching for the self within these abstract perspectives. We suggest a different starting point and strategy for developing models of a self which is more contextualized within the realm of action. The idea is that within a more contextualized framework one is able to formulate a theory that is ‘closer to the ground’ and less abstract.

Insights developed in certain neuropsychological studies suggest a way to define the deficiency of the above-mentioned approaches, and to project a path that would

[1] Sartre (1957) views the notion of the ego as a defence mechanism that protects us from the anxiety associated with an authentic realization of our absolute freedom. Recently, Varela, Thompson and Rosch (1991), integrating a Buddhist psychology with a cognitive science approach, conclude that ‘the ego-self that everyone clings to and holds most dear,’ doesn’t really exist, and that this fact can be ‘profoundly transformative’ (pp. 80–1). However, as the authors acknowledge, whether *this* kind of a self exists or not, does not rule out the existence of other kinds of self. For the concept of ‘centre of narrative gravity’ see Dennett (1992).

lead to a more comprehensive and less abstract model of the self. Moreover, for this more comprehensive model, considerations about agency and ethical action are most pertinent. We are led to a perspective that takes ethics (in the most general sense of ‘having to do with how one lives one’s life’) as a suitable starting point for working out an understanding of the notion of self. In contrast, in most traditional philosophical approaches to the question of personal identity, the starting point is purely epistemological or metaphysical, with the result that only in the end and as a seeming after-thought does one try to sort out the implications for the ethical realm of action.²

On this score, Galen Strawson’s (1997) model of the self is not unconventional within the discourse of philosophy. He sets out to develop a phenomenologically informed metaphysical (specifically, for Strawson, materialist) conception of the self.³ This approach leads him to a conception of the self as a relatively short-lived mental thing, a temporary subject of experience. Along the way he mentions, but then excludes from the essential aspects of self, the ideas that the self can ordinarily be thought to have character or personality, and to be conceived as an agent. He offers only a short discussion of personality in which he takes the state of depersonalization to be the essential aspect of a mental self. This he describes as a ‘bare locus of consciousness . . . void of personality, stripped of particularity of character, a mere (cognitive) point of view’, which one might experience as ‘the result of exhaustion or solitude, abstract thought or a hot bath’. Strawson contends that a sense of self such as this survives depersonalization, and since even in normal circumstances personality is something that goes unnoticed and undetected as an object of experience, personality and character are accessories, not essentials to the sense of self. In Strawson, then, no less than in the philosophical tradition that stretches from Locke to Parfit (see, e.g., Parfit, 1984), the ethical dimensions of self are usually explored only in terms of what implications or consequences an already worked-out conception of self may hold in such respects.

We have two final introductory remarks. First, we want to be clear that although this paper is centrally concerned with the nature of the self, there is a necessarily related issue that we address, namely, the question of access to the self, and whether there can be certain forms of self-consciousness that are not abstractions from contextualized activities. The promise of a sound basis for the development of a theoretical conception of a contextualized self is only good if in fact there are reliable forms of contextualized self-consciousness, since the primary method for getting a grasp on the self is through first-person self-experience. Beyond this, however, the question of access is essentially linked with the question of the nature of the self. Consider an animal that has no experience or awareness of its actions. It has no access to something that we would call self. We are inclined to say that such an animal *has* no self. The question of self or personal identity is an issue only for an animal which has some access to itself within the context of its own behaviour; *access to ‘itself’* actually helps to make possible *the existence of ‘its self’*. Access (self-consciousness) is constitutive of self.

-
- [2] Carol Rovane’s recent book, *The Bounds of Agency* (1998) makes a similar point. Her analysis, which ends up defending the notion of psychological continuity, leads in a very different direction from the one taken here however.
- [3] As several commentators have already pointed out, Strawson’s phenomenology is not the methodological phenomenology developed by Husserl and his followers, but more of an informal introspection. See Sheets-Johnstone (1999), and Zahavi and Parnas (1998).

Second, we wish to be clear that in sketching an approach to a conception of a self in contextualized action, we do not assume that there is only one kind of self or that an explanation of the contextualized self will be an explanation of every sense of self. Other approaches, such as the Meadian analysis of a socially constituted self, or the notion of an autobiographical self, can reveal important and valid conceptions of self.

Intentional Attitudes and Contextualized Action

Central to our proposal are two concepts: the *intentional attitude* of a person and *contextualized action*. The intentional attitude of a person consists of the content of their current purposive intentions which is itself a function of their attentional focus. It is inferred from the performance of perceptual, motor and linguistic activities, from phenomenological report, and also from the task demands and external situation. The total context of an action is made up of the intentional attitude and the external context, which we see as largely but not totally interdependent. Intentional attitude and contextualized action will be more fully specified below. Different kinds of contextualization can be distinguished by considering the variety of ways in which the effects of brain damage are affected by or mitigated by factors most easily described on the personal level, especially in terms of the intentionality of the behaviour involved. However, while we make broad categorizations of intentional attitudes and contextualizations in regard to the data discussed below, we suppose that they vary continuously and are a complex mix.

In certain approaches to rehabilitation of function after brain damage attempts are made to elicit the behaviour that is impaired or inaccessible.⁴ Apart from the rehabilitative motives for this, if it is possible to elicit otherwise inaccessible behaviour or function, that which is effective in eliciting it is of great and obvious theoretical relevance. Consider first the implications of Leontiev and Zaporozhets' (1960) research on rehabilitation of impairments in hand use. They showed that in some cases hand movements can be more effectively rehabilitated by having the patient perform the impaired behaviour in the context of meaningful activity than in the exercise of isolated movement. In patients suffering from ideomotor apraxia, who are otherwise unimpaired in perception, comprehension, or motor performance, actions that cannot be produced on request or by imitation, can be produced or improved when they are performed in the course of normal activities that include such actions. Marcel (1992) found this same phenomenon in experiments with motor-impaired neurological patients who showed characteristics of ideomotor apraxia in manual function but were not classified as ideomotor apraxics since their motor impairments were identified as relatively peripheral (rather than of central origin). These patients showed a significant improvement in various aspects of motor control and fluency in impaired behaviours when performed as meaningful actions, over their performance when elicited as decontextualized behaviours. More significantly, in almost all of these cases even further improvement was found when nominally the same movements

[4] For example, in the therapeutic procedure called 'deblocking' or 'stimulation' technique (Weigl, 1961; Weigl and Kreindler, 1960). The assumption made in the stimulation approach to therapy is that the impaired function is not completely lost, but rather is inaccessible. The attempt is made to find an appropriate method of elicitation so that the patient might be able to recover, rather than relearn, the function in question.

were performed in a situation, usually a social situation, in which the movements constituted actions with personal and culturally derived signification.

For example, a woman who had difficulty in grasping, in lifting, and in motor fluency when asked to lift a cylinder of the weight and size of a glass of liquid and to move it toward or away from herself, showed clear improvement in her performance when spontaneously drinking during a meal. This same woman was even more proficient, almost normal, in the very same movements when serving mugs of tea to guests in her home, although not when clearing up the mugs. A second patient who had coordination, timing and sequencing problems in finger control, found it difficult to copy letter-like figures, but improved when writing words to dictation, and performed best when writing her plans and achievements in her diary.

On the basis of such observations and of follow-up experiments, Marcel (1992) distinguished three levels of performance, the baseline level plus the two levels of improvement. These levels of performance turned out to be associated with two things: the degree of semantic or pragmatic contextualization,⁵ and social or self signification. In all but one of the patients examined by Marcel, differences in the nature of intention strongly correlated with performance differences. Marcel labelled the performance levels, simply, Levels 1, 2 and 3. Level 1 (worst) performance was obtained in situations in which the patient was instructed to carry out a disembedded, meaningless or purely procedural action (often the case in psychological or neurological experimentation or examination). Level 2 performance was obtained in contextualized actions or activities that involved a pragmatic significance or purpose relative to the patient (in the above examples, drinking at a meal, washing dishes, writing to dictation). Level 3 (best) performance was obtained in actions that were personally significant or that derived their signification from the social and cultural system (in the above examples, serving tea to friends, writing in a diary; in other cases, cutting bread or meat for others at a family meal, dealing cards at a real game). Level 3 performance, involving improvement to normal or near-normal performance, appeared to be associated with intentions which implicitly refer to self, in those cases in which cultural practice usually assigns such symbolic signification. For example, grasping mugs of tea and giving them to guests is underlain by the intention to 'offer hospitality', a cultural practice which contributes to the constitution of the social competence and self-esteem of the agent. Following such changes in intention, a specific motor performance can move from impossible, or near impossible (Level 1), to possible or improved (Level 2), to relatively fluent (Level 3). Actions do not need to be performed in the presence of others to have a social signification. However many of them were so performed. One might think that performance of such actions in the presence of others would lead to greater 'self-consciousness' which would *disrupt* the fluency of action. In experiments, Marcel found that this did often occur. What is important is that in real social situations, after initiation of the activity, the opposite (greater fluency) occurred.

In the cases examined by Marcel, different performance levels appeared to correspond to differences between three kinds of intentions.

[5] It will be clear from the following analysis that 'semantic or pragmatic' refers to a situation in which there is a meaningful intention on the part of the agent, or in which meaningful action is possible.

- Intentions that are relatively *abstract or decontextualized* (i.e. that have no purpose beyond the task itself or compliance).
- Intentions that are *pragmatically contextualized* (e.g. those involving purposive behaviour, or the exercise of an already known intention).
- Intentions that are *socially contextualized* (e.g. those involving self-reference or other persons — where the self is socially embedded by action and where actions both express and signify competence in a social and emotional role).⁶

We can enlarge on these distinctions by first defining them in terms of the contextualization of actions which correspond to them, and then by characterizing their respective attentional foci.

An *abstract or decontextualized* action is one which is detached from what would ordinarily be considered a significant context or where the person has no normal or good reason for doing what is asked other than voluntary compliance. A good example can be found in an experimental situation in which the experimenter asks the subject to perform a simple movement, e.g., touching one's own nose. A *pragmatically contextualized* action is one that is relatively more informed by a meaningful structure (relative, of course, to the individual subject) and is performed in the course of a natural activity whose purpose arises from personal projects and concerns. These two kinds of action closely align with Kurt Goldstein's (1940) distinction between *abstract* and *concrete* behaviour. Concrete behaviour, corresponding to pragmatically contextualized action, is behaviour that a subject performs and has a reason to perform in a situation that is closer to real-life (e.g. scratching, or swatting a fly from one's own nose, or showing someone else where a mark is, in contrast to the behaviour of touching one's nose on command, as in the clinical test).⁷ A *socially contextualized* action is not only one where the agent is embedded in a social context, but where the action has meanings defined by cultural categorizations ('gestures') and represent states of the self in regard to others. It is rare that pragmatically contextualized actions are not also socially contextualized. However, this does not mitigate the distinction.

It is important to note that a specific kind of behaviour or movement cannot be categorized *per se* as having any one kind of intentional or contextualization status exclusively. The distinctions between these actions are relative to the agent's intentions, and to different reasons for action or the relative absence of reason (e.g. for no reason other than compliance in the case of an experiment). To point, or touch, or scratch: these are not *intrinsically* abstract or concrete behaviours; their status in this regard depends on intention. For example, a movement of one's hand to a specific

[6] These three kinds of intention can be exactly exemplified by three kinds of utterance, all of which have the 'objectively' identical external phonological form.

(a) You agree to read aloud a list of pronounceable letter-strings that are orthographically non-words, one of which is PHYRE.

(b) You answer the question 'what results if you set light to dry wood?' with the word 'fire'.

(c) You see that an office has caught light and is burning and run out shouting the warning 'fire!'

The manner of the utterance of identical phonology would differ, but crucially it would do so because the content of the intention would differ, and this varies in accord with the situation.

[7] Like the patients discussed above, Goldstein's neuropathological patients often failed on tasks defined on an abstract level. Patients would be unable to perform abstract movement but could perform the mechanically equivalent movement in more concrete situations.

part of one's own body may be (nominally) mechanically identical in all three levels of action described above.⁸

The intentional attitudes of the distinctions sketched above are characterized by differences in their attentional focus. The notion of an abstract or decontextualized action indicates one characterized not only by a lack of contextualization but also by some degree of relatively high-level cognition or conscious attention directed to the particular behaviour itself (e.g. understanding instructions and deliberately translating them into motor actions that one monitors). When, for example, I am told to walk a straight line as part of a neurological examination, I am more conscious of my movement than if I am simply walking to open a door for a friend. The notion of a pragmatically or socially contextualized action indicates one with higher degrees of intentional content (one that is purposive and involves intentions that go beyond that behaviour itself) and a relatively low-level of attention targeted to the particular behaviour in question. So, for example, in walking to open the door for a friend I will not be thinking of (or possibly not even be aware of) how I am moving my legs; one might say that my walking is an unthinking (close-to-automatic) aspect of the larger intentional action. It is also possible that the very same motor functions which formed the conscious theme of the subject's activities in the abstract situation are, in a more contextualized setting, incorporated into and become an intrinsic part of a larger task that bestows meaning on the particular movement. The task itself may not be particularly meaningful (for example, washing dishes within an experimental or therapeutic situation); but relative to the task a particular hand movement takes on meaning and is generated by that meaning.

Of course, in many social situations what matters to the agent is indeed the *manner* of the action, e.g., walking elegantly or nonchalantly. We suggest that disruption of the intended behaviour in such cases is due to the behaviour being the *explicit focus of consciousness* rather than an *implicit aspect of the intention*. For example, it is important for the actor, dancer or athlete, once having worked on such aspects of movement, to demote them from the focus of conscious intention in performance. The main point here is the difference in intentional attitude when the behaviour in question is the explicit focus of conscious attention as opposed to when it is an implicit aspect of the intention.

One can find numerous examples of behaviour on these different levels across many different behavioural domains: perceptual, cognitive, linguistic, motor, and emotional.⁹ Marcel and his colleagues presented auditorily, single, three-syllable nonwords (e.g. 'miladu') to a patient with conduction aphasia (a form of aphasia that involves difficulty with repetition, especially of nonsense words). The patient was asked to repeat each of them. She was able, at best, to repeat single syllables, sometimes combining the consonant and vowel of different syllables, sometimes making phonemic paraphasias. When the task was redefined and the requirement of repetition was an *implicit* requirement of a request simply to say which one of three auditorily presented nonsense words she preferred (using the same stimulus set as before), now the patient succeeded in 28 out of 30 trials in properly remembering and pronouncing one of the three-syllable nonwords. This task implicitly required repetition, but that

[8] Of course there are what O'Shaughnessy (1980) calls 'sub-intentional actions', such as tongue movements while writing, that are irrelevant to the main activity and are voluntary, but 'unaware'.

[9] Goldstein (1940) presents numerous examples.

aspect was embedded in the intention of telling which word she preferred, a task that included an aspect of personal intentionality. Repetition, in this case, was made possible by making it an implicit rather than explicit part of the agent's intentions.¹⁰ Of course the total situation had now changed, because of what was being asked of the patient; and thus the role of the questioner was different (an enquirer after personal preference rather than a tester of competence in a task known to be frustratingly difficult). This in turn changes the intentional attitude, from attempting to repeat to expression of preference.

That the meaning content of an intention according to which the action is performed has an effect on performance, can be seen in a study by Brouchon *et al.* (1986). In both optic ataxic patients and in normal subjects, perceptual-motor coordination in what is mechanically the same movement is qualitatively influenced by whether the instruction is 'to reach for and touch' or 'to indicate' an object. Movements performed under different descriptions or instructions entail differences in the subject's intention; the second instruction insofar as it requires deixis (pointing) implicates communicative reference to another person.

Another example is both clear and dramatic. It involves congenitally deaf users of American Sign Language (ASL) who suffer from hemispatial neglect following stroke. In cases of left hemispatial neglect (Jeannerod, 1987) subjects entirely ignore their left perceptual and/or motor field, or the left side of their body. Such patients do not refer to things in their left hemispace and have a tendency to fail to perceive or attend to what is to their left. They obviously experience difficulty with spatial tasks that are in or that refer to left hemispace; so, for example, they may be unable to map out a room or properly describe the layout of objects in space. It is important to note that some of the syntactic and discursive aspects of ASL make use of space. For example, when entities are introduced in discourse they are each assigned a spatial location around the signer which is returned to for anaphoric (especially pronominal) reference. The same ASL users with acquired hemispatial neglect who ignore their left and have profound difficulties with tasks in their left perceptual field and left peripersonal space have no problem or pathology in discursive or syntactic use of left hemispace when involved in socially contextualized communicative acts of signing (Poizner *et al.*, 1987). This suggests that it is what the space represents, within an intentional context, that matters. In the conversational situation peripersonal space, rather than representing only itself, represents the reference and pragmatics of discourse for the subject acting in a culturally determined sign system, just as phonology does for the normal speaker.

Changes in the content and organization of the agent's intentions result in changes in the performance level of action. In many of the cases cited, even in normal subjects, behaviour tends to break down and become disintegrated in decontextualized tasks. In such tasks the opportunity has been reduced for the content of the intention to go beyond the immediate task. In pragmatically and socially contextualized tasks

[10] See Marcel (1992) for summary of this study. The patient's performance following a redefinition of task provokes the question of whether the patient's short-term auditory-verbal memory was impaired or whether her difficulty involved access to and use of her memory. Goldstein describes a similar case involving memory. 'The patient has the material in his memory, but he is unable to use it freely; he can use it only in connection with a definite concrete situation, to which it must seem to him to belong' (1940, p. 51).

behaviour tends to be more integrated, presumably because the agent's intentions encompass more than the immediate action itself, their attentional focus goes beyond it, and its significance is part of the larger projects and concerns of the person.

These generalizations do not hold in all cases, however. There are many pathologies (for example, some instances of aphasia and associative agnosia) where patients perform worse at the semantic or pragmatic levels and better at the lower or more abstract level — some aphasic patients may be able to apprehend sensory and lexical characteristics of words, but not their meaning; many agnosic patients can perceive sensory aspects of objects but fail to know what the object is or how to use it even when in context (see Marcel, 1983). It is difficult, however, to find corresponding clearcut instances in normal subjects for two reasons. First, normal people have the ability to behave in a decontextualized manner, even if less well; and can bestow meaning on it. Second, conscious attention is normally directed to the highest semantic or functionally useful level of behaviour — in perception, at the level of objects and action, in language, at the level of communicative pragmatics. Difficult, but not impossible. In the midst of a highly contextualized social situation, the normal person may suddenly become overly self-conscious of his/her own movements or speech and suffer a corresponding performance decline. When a poorly skilled reader devotes more attention to coding processes, he/she may lose consciousness of meaning (see Marcel, 1983). These examples, however, also indicate how the availability and execution of behaviour is influenced by the description or the intention under which such behaviour is generated. Personal goals and motives, as well as the social pragmatics of the situation (always defined relative to the subject), can clearly provide the energizing mechanism which transforms intentions into action. For our purposes here, the distinctions between performance levels are useful simply for indicating and providing evidence for the distinctions between intentional attitudes, even if performance levels are not strictly correlated with specific kinds of such attitudes in all cases.¹¹

At this point we can fill out our definition of intentional attitude and add some comments. The intentional attitude is characterized not only by its level of contextualization but by its attentional focus (e.g. on the action itself versus its function or signification: lifting a glass, drinking, offering hospitality) and correspondingly the breadth of that focus (whether it is restricted to the action *per se* or goes beyond it temporally or socially). Attentional focus can be more radically divergent: on the self or the world, inner or outer directed. In such cases it is more evident that its breadth is relatively independent of its object or directionality. It is worth emphasizing that the increased breadth of attentional focus, which corresponds to contextualization, does not imply that all such content is explicit in the person's conscious experience. When the act of passing mugs of tea is embedded in the larger project of hospitality to guests, one's intentions extend to that project, but it is rare that one is immediately aware of all of this or of its significations.

The intentional attitude is also characterized by the *mode* of attention. In an *observational mode* one is detached from that which is the focus of attention; in a *non-observational mode* one is immersed in it. This corresponds to being in a state of reflexive or nonreflexive consciousness respectively. Furthermore, while there is a

[11] The relation between performance level and intentional attitude as defined above, will vary across individuals and across different pathologies.

correlation between the external situation and the intentional attitude, they are independent. In one and the same situation a person can change their attitude. Indeed this is just what is learnt in certain meditation and therapeutic techniques. For example, in depression the cycle of negative thoughts about the self and the inability to take action which maintain and increase the mood can be broken insofar as the person is able to learn to attend to the world and projects beyond themselves (see Teasdale, in press).

The importance for the concepts of self and self-consciousness of the foregoing discussion will become evident in the context of the following examination of the unity of consciousness.

Unity and Abstractness

The question of the unity of consciousness is one that is often raised in philosophical and psychological investigations into the nature of the self. Our intention is not to address the difficult problem of the unity of consciousness, but to examine the ways in which philosophers and psychologists attempt to address this problem. We want to suggest that various answers to this question are shaped by the kinds of behavioural situations in which philosophers and scientists are inclined to conduct their investigation.

Nagel, for example, takes issue with an assumption which underpins the idea of a mentalistic self, and which thus remains operative in many philosophical analyses of the self. The suspect assumption is what we might call the *assumption of the unity of reflexive consciousness*: that ‘a single mind has sufficiently immediate access to its conscious states so that, for elements of experience or other mental events occurring simultaneously or in close temporal proximity, the mind which is their subject can also experience the simpler *relations* between them if it attends to the matter’ (Nagel, 1975, p. 239). Through an examination of cases of split-brain patients Nagel is led to a certain scepticism: ‘I believe that consideration of these very unusual cases should cause us to be skeptical about the concept of a single subject of consciousness as it applies to ourselves’ and, by extension, to others (p. 242). Rather than thinking in categories that involve the unity of consciousness, he suggests, we need to think in terms of behavioural functions that are normally integrated but that can be dissociated in either experimental situations or cases of brain damage.

If one adopts Nagel’s perspective, if, for example, one examines dissociative states of consciousness, it seems quite reasonable to question the notion of a unitary self (Marcel, 1993). As Nagel points out, it is not unusual to assume that the unity of self is based on the unity of consciousness. But then, on close inspection, the unity of consciousness does not seem to hold up to philosophical reflection. One need only think of Hume’s famous and failed introspective attempt to locate the self in consciousness, and his conclusion that what we call the self is nothing but a bundle of impressions; or of Strawson (1997), who presents a similar although less bundled example. One possible response to such conclusions is to ask about the agent involved in doing the reflecting, so that, even if Hume consistently finds a bundle of different impressions, he does so from a consistently unified perspective (see Gallagher, 1992). This idea, again, involves the assumption of the unity of reflexive consciousness. The argument would be that even if phenomenal consciousness is not unified, at least our (second-

order) access to it is, and that unity provides, in part, for the veracity of a more detached reflection.¹²

The same assumption, that the subject's reflexive consciousness is unified, operates in many psychological studies. Psychological experiments are often set up on the assumption that the test subject has a unified consciousness.¹³ This assumption is implicit in the fact that the mode of report (the subject's mode of response) is not expected to make any difference in the content of what subjects actually report as their experience. This is clearly a Cartesian assumption: subjects have direct access to their experience (e.g. either they see or they don't see a certain visual stimulus) and the way they report on that experience is immaterial to what that experience is. Consider, however, the following experiments which suggest disunity in reflexive consciousness.

In a set of experiments that began as studies on a subject with blindsight, Marcel (1993) was led to test normally sighted subjects. In one experiment 10 normally sighted subjects were asked to discriminate on each trial between the presence and absence of a luminance increment in a target light, using a threshold increment value which yielded for each subject 50% correct guesses and 25–30% reports of 'definitely seen'. On each trial subjects responded simultaneously in three response modes: by eye-blink, by finger-press on a button, by oral 'yes/no'. In different blocks of trials they were asked either to *report their experience* of a luminance change or to *guess presence/absence* of a change. In both cases they were asked to respond 'as fast as possible but as accurately as possible'. Latencies of the responses indicated that they were not reflexes. The results for guessing showed very little difference between the response modes. However the dramatic finding was that in the condition requiring report of experience, on identical trials the responses dissociated across different modes. For example, people would simultaneously indicate 'yes' with their eye-blink but 'no' with the finger-press, or 'yes' with both eye and finger but 'no' with the oral response, and the reverse also occurred. Overall accuracy of report (correct vs. false positives) differed significantly between the three modes; it was best for eye-blink and worst for oral report. Although subjects correctly remembered the instructions, when reports were made fast they did not realize (even when questioned immediately

[12] Even detached reflection has experiential and, in that sense, phenomenal content (and this means that it can become the object of a further reflection). Insofar as reflective consciousness takes conscious experience as its object, as in the kind of reflective introspection often found in philosophical analyses of the self, it is twice removed from the conscious experience that is its object, and depends on the more immediate reflexive access we have to experience. This kind of reflective stance is taken by Edmund Husserl, for example, in his complex phenomenological analysis of the temporal stream of consciousness. Despite his trust in reflection, however, Husserl notes:

We must therefore distinguish: the pre[reflective] being of experiences, their being before we have turned towards them in reflection, and their being as [reflected] phenomena. When we turn towards the experience attentively and grasp it, it takes on a new mode of being; it becomes 'differentiated', 'singled out'. And this differentiating is precisely nothing other than the grasping [of the experience]; and the differentiatedness is nothing other than the being-grasped, being the object of our turning-towards' (Husserl 1991, p. 132).

Husserl means that the reflective act lifts the reflected part of consciousness out of the stream of consciousness, freezes it when in fact that freezing distorts its nature. So the reflecting is just this differentiating of the reflected part from the stream. For further discussion, see Gallagher (1998).

[13] Even the experimental examination of split-brain patients starts with this assumption since instructions are given to the subject as if he or she were a single individual.

after a trial) that there was a discrepancy between the response modes of their reports. This discrepancy did not occur in the guessing conditions.

In another experiment, again with 10 normally sighted subjects and again with either guessing or report, each response mode was tested separately, i.e. on a block of trials subjects responded with only one kind of response on each trial. Also in separate blocks of trials subjects either had to respond as fast as possible or had to wait for a signal, delayed by 2 or 8 seconds, before responding. It was found that relative accuracy was the same for each response mode as when they had been performed simultaneously. Again it was found that in all response modes guessing was more accurate than report, and that accuracy did not differ between response modes for guessing, but did so for reporting. In neither experiment was the non-difference for guessing due to performance being at ceiling, i.e. as high as logically possible. This second experiment shows that the discrepancy between report modes was not due simply to having to make three responses on each trial, nor to the order of or delay in such reports.

The importance of this is that traditionally, in psychological experiments on non-conscious perception and blindsight, greater accuracy of guessing than reports is interpreted as reflecting the presence of nonconscious perception versus conscious perception. Nonconscious information is taken to guide or have a causal effect on guessing, while having a very much smaller role in affecting the rational descriptions, in report, of the subject's conscious perceptual experience.

How should we interpret these data? First, conscious detectability of experience appears to differ across different modes of report. Subjects on the same trial are reporting with a button press that they see a light and reporting orally that they do not see a light, or vice versa. The second experiment shows that this contradiction does not depend on the requirement of making simultaneous reports; that is, it is not a property of the requirement to make more than one report on a single trial. If reports are indeed reports of experience¹⁴ then it appears that an experience is not independent of report, but depends on the mode of the report. This may mean that the nature of an intended report influences experience,¹⁵ but since in the experiment involving simultaneous reports all three report modes were intended, this cannot be a complete explanation. An alternative and preferred hypothesis is that different ways of reporting have differential access to an experience. This calls into question the existence of a unitary reflexive consciousness, or a unitary subject of experience responsible for report.¹⁶ It suggests not only a distinction between phenomenal experience (the sensed experience of the light) and reflexive consciousness (the awareness that we

[14] There are several good reasons to believe that they are reports of experience. First, spontaneous latencies are too long to be reflexes. Second, subjects, when asked, insist that their responses are 'reports'. If people deny awareness of such stimuli, they are usually unwilling to make voluntary responses (see Marcel, 1988). Finally, there are significant differences in results between responses in the case of reports of experience and guesses in the absence of conscious perception.

[15] Bisiach, Berti and Vallar (1985) suggest that the nature of the response used to report a visual experience can alter the memory of the experience in as little as two seconds.

[16] Similar dissociation in reflective consciousness was found by Cumming (1972). Subjects were shown letters briefly and laterally masked, and were asked to indicate the presence of a specified target letter by pressing a key. Asked to respond quickly, subjects 'tended to respond to the presence of a target letter with a fast press of the "target-present" key, and then, a moment later, to verbally apologize for having made an error, when in fact they had been correct'.

are experiencing something, which becomes the basis for the report), but disunity in the latter. The apparent contradictions revealed in the experiments (e.g. a subject's reports that he is both experiencing the light and not experiencing the light) are resolved on the supposition that the experience is differentially available to divisions of reflexive consciousness.

The supposition of disunity in reflexive consciousness is reinforced in cases of anosognosia (when patients demonstrate a lack of awareness of their deficits or impairments). In some cases, patients seem to demonstrate both awareness and non-awareness of their neurological problems. For example, a patient may verbally admit or complain of his hemiplegia (unilateral paralysis), and at the same time try to act as if he were normal. On the other hand, some hemiplegic patients appear unaware of their deficit, and deny it if questioned, but never attempt to initiate activities requiring coordinated use of limbs on both sides (Bisiach and Geminiani, 1991).

In another set of observations (Marcel, 1993; Marcel *et al.*, under review) sixty-four hemiplegic patients who were anosognosic for their plegia were asked to rate (1) their own ability on activities that involve coordinated activities of limbs on both sides (e.g. tying a knot, clapping) and (2) the questioner's ability to do such tasks *if the questioner were in the same condition as the patient*. Half of the patients were asked to rate (1) first, and (2) later; the other half were asked to rate these in reverse order. In such instances, many patients rate their own ability as perfect but claim that the questioner, if in their own (the patient's) condition, would not be able to perform the task, 'because you would need both hands'. In each case (1 and 2) a reflective stance is taken, but in (1) patients fail to recognize their deficit while in (2) they affirm it. Patients failed to notice the inconsistency between their responses. Some of these self-contradicting patients showed another dissociation of awareness involving self-knowledge. When asked whether each arm was at all weak they casually denied any weakness in their plegic arm. However, when asked immediately after, in a confidential manner, 'Is this arm ever naughty? Does it ever not do what you want?', several patients, although bemused at this question when referred to their nonplegic arm, answered affirmatively when referred to their plegic arm. (One patient said 'Oh yes! In fact if it doesn't do what I want, I'm going to hit it.') It seems that patients' awareness of their state and abilities is split and contradictory, and depends upon the kind of personal stance that they take.

Because reflexive consciousness is thought to have a close connection with selfhood (Kihlstrom, 1993; Marcel, 1993; Nagel, 1975), such considerations lead again to scepticism about the concept of a single subject of consciousness. If reflexive conscious states can be dissociated, then there is little reason to assume a unified self, or even a single self in one body or brain.

Nagel (1986) moves beyond the assumption of the unity of reflexive consciousness and furthers the discussion by considering an alternative hypothesis to mentalism, namely, the equation of self and brain. Not unlike Strawson's stance on materialism, he embraces the 'dual aspect theory' of self, and maintains that brain states are both physical and mental. For him, this means that the referent of the psychological subject is actually the brain, something 'which is the persisting locus of mental states and activities and the vehicle for carrying forward familiar psychological continuities when they occur.' Nagel continues:

I could lose everything but my functioning brain and still be me. . . . If my brain meets these conditions then the core of the self — what is essential to my existence — is my functioning brain. . . . I am not just my brain. . . . But the brain is the only part of me whose destruction I could not possibly survive. The brain, but not the rest of the animal, is essential to the self (1986, p. 40).

One may thus be tempted to push aside the perplexities involved in dissociations of reflexive consciousness and the problems of psychological discontinuity, the subject of seemingly interminable philosophical discussions, by appealing to reductionist neuroscience.¹⁷ Yet, reductionist or functionalist interpretations of neurological evidence do not appear to settle the philosophical problems of the self, since the diversity of such interpretations make it possible to develop various and incommensurable models of the self that simply mirror the philosophical theories.¹⁸ If the framework of inquiry is set by the usual kinds of questions posed in philosophical investigation or at the level of psychological or neurological experimentation, then the answers that seem most natural are those that are framed in terms of consciousness and/or the brain. Here again we find the traditional problems and paradoxes that keep the issue completely unsettled, and that make up the majority of the philosophical and psychological history of these concepts. What we need to be clear about, however, is that in all of these approaches the self is being sought from an *abstract and detached* behavioural perspective.

We have characterized abstract behavioural situations in terms of a decontextualization. It is also possible that a subject may actually be reduced to abstract or detached behaviour by various pathologies or in certain limit situations (involving fatigue, illness, and so on). We want to suggest that philosophical methods which involve introspective self-consciousness or the hyperreflective analysis of behaviour may also place the reflecting agent in an abstract behavioural situation or a detached stance. Reflective self-consciousness (the unity of which is itself in question) involves something like artificially (and sometimes experimentally) pulling back from particular contextualized activities and posing a question from an abstract or detached point of view. Not only questions such as ‘What is the I?’ or ‘What is the self?’ or ‘How can I account for the unity of consciousness?’ but also the practice of reflective introspection itself shift and redefine our intentions away from what may have been an active engagement with the world. As in some experimental situations, such questions involve shifting the focus of attention away from purposive activity involving meaningful objects and other persons, to our own movements and modes of consciousness. Such hyper- reflection is a third-order cognitive activity, once

[17] This temptation is often accompanied by a certain discounting of the difference between the discourse of neuroscience and the personal-level discourse involved in the description of self and self-consciousness. There are many objections to this move. One is Dennett’s (1969) — that different entities are involved at different levels of discourse. Another is simply that when someone refers to their self, they are not referring to their brain, with which they have no acquaintance. For our purposes here we forego any objections that are rightly made to this conflation of discourses.

[18] Contrast, for example, the details of Dennett’s (1991) abstract centre of narrative gravity, Damasio’s concept of a ‘neural self’ which, not unlike Galen Strawson’s model, is, at each moment, ‘constructed, from the ground up’ (1994, p. 240), and the account provided by Ramachandran and Hirstein (1997) which regards the self as a delusion. Disagreements among these accounts involve both philosophical and neurological details.

removed from reflexive consciousness, and twice removed from phenomenal experience or the behavioural level at which we find contextualized action.

Donaldson (1978) has pointed out that such academic and analytic attitudes are the goal of much of Western schooling and involve the ability to detach one's consciousness from determination by worldly and pragmatic contexts. Such attitudes are often difficult to learn because they involve addressing specific aspects of pre-semantic levels of representation in perception and action, when we are normally inclined to operate at the semantically or pragmatically contextualized level. For example, in order to learn to read (new words) in an alphabetic orthography, one has to attend to auditory words at the level of the sequence of phonemes and to desist from attending to their wholes and their meaning.

It is in this kind of detached, abstract, analytic attitude that, as philosophers or scientists, we start to search for and account for the self. To ask, for example, 'What is the self *per se*, or essentially, or *in itself*?' is to ask 'What is the self, apart from or outside of any particular context?' Philosophical and scientific questioning itself may be quite contextualized in social and professional settings. But the very nature of the questions (for example, 'What makes me or someone else an identical self across any number of particular contexts and behaviours?') and the preferred methods for answering them, lead us to seek answers that are abstracted and that exist in abstract decontextualized behavioural situations. In such an approach, one looks for something relatively general — the subject of experience, independent of any particular experience, and thus independent of any particular context.

Methodological frameworks in philosophy and the sciences lead us to believe that we are working out solutions on a relatively basic level when we do this kind of analysis. It is clearly a mistake, however, to think of the more abstract behavioural level as more basic than contextualized behavioural levels. The experimental data cited in the first section (i.e., the fact that, at least in some cases, performance of abstract behaviour is frequently impaired first and to a higher degree while performance remains more intact in pragmatically and socially contextualized action) suggests that abstract behaviour may not be more basic than contextualized behaviour. Socially contextualized behaviour is not built out of abstract behavioural components; rather, abstract 'components' are simply abstractions of more contextualized behaviours. We do not intend this as a developmental statement — that is, our claim is not that more complex social behaviour is not built up from more simple behavioural components as the individual matures, etc. What we want to say is that complex social behaviour is more than the sum of its parts (indeed it is not clear that it really has independent parts), and that meaning (or meaningful action) cannot be analysed in purely mechanistic or atomistic terms.

It is important to note that dissociations such as those in reflexive consciousness shown by the experiments cited above, are not usually observed by either the subject (since one can logically never directly experience more than one consciousness or self) or even by others. What we come to consider to be the self from the perspective of abstract behavioural situations like those produced in philosophical reflection or in much scientific experimentation, is no more than a detached snapshot of a self that quite possibly functions in a more integrated manner in socially contextualized situations.

Our claim is not that reflexive awareness and a more detached reflective consciousness¹⁹ are not involved in the workings of the self, but that by focusing on detached or decontextualized states philosophers like Hume, Nagel and Strawson, as well as psychologists and neuroscientists, are searching for the self in the wrong way, or certainly in a way that will give a partial and distorted picture. If the self cannot be accounted for purely in terms of reflexive consciousness or purely in terms of brain functions (both, of course, necessary conditions), neither can it be adequately accounted for by a dual-aspect theory that would locate it in a mix of brain and reflexive consciousness, if that theory is developed in a framework that is built around abstract behavioural situations. The answer, however, is not to abandon the notion of a self, but to start the search at a different level.

The Self in Personal and Social Pragmatics

We have argued that in general, and for the most part, philosophical reflection operates in such a way, and psychological and neurological experiments are designed in such a way, as to limit the subject to abstract intentional attitudes or to a detached stance. And it is within the limits of such situations that philosophers and psychologists pose and answer abstract questions about the self, the unity of consciousness, and so forth. It is to be noticed that this limitation holds for both first- and third-person approaches.

Is it possible to develop an approach to the self that would take its bearing within pragmatically and socially contextualized action, i.e., where the subject is in a non-reflective intentional attitude? What does the self look like in such behavioural situations? What kind of access do we have to the self in such cases? If we take a third-person perspective we easily discover, not a unitary phenomenon, but a self with multiple but relatively integrated aspects. That is, as social psychology often suggests, the subject (i.e. the actor) plays different social roles within different social contexts. As an academic or scientist my socially contextualized behaviour is typically circumscribed so that I represent myself in a certain way. As a family member, citizen, religious practitioner, sports enthusiast, my various activities are characterized in relatively different roles and vocabularies, and so on. So in my various activities, I am many different selves to many different social groups. Following this approach it seems that even in socially contextualized situations, we can discern multiple selves, albeit in a somewhat different way from Nagel or Strawson. It is not uncommon, however, to think of these different roles or aspects as being integrated in some relatively rough way. The concept of ‘character’ or ‘person’ is sometimes used to indicate that an agent has or manifests some relatively stable and consistent characteristics across all of these roles. This relatively integrated agent in some way constitutes what is ordinarily called the self.²⁰

[19] By reflexive consciousness we mean a second-order awareness directed at and capable of reporting phenomenal experience; reflective consciousness is a detached, offline or introspective consciousness.

[20] Kathleen Wilkes (1988) argues in this way. Pragmatic support of this view can be found in legal and political systems that treat one agent as responsible for actions, no matter under what role description the action has been perpetrated (unless the agent is deemed to have been in a fugue state). In short, one can sort this issue out with relative ease on third-person approaches since there are acceptable and practical ways to adjudicate alternative descriptions (multiple selves vs. unitary agent) of socially contextualized behaviour.

The approach we outline below is not inconsistent with such third-person views. Our question, however, concerns the possibility of gaining *first-person* access to the self in socially contextualized behavioural situations. For our purposes here, we leave the third-person framework aside in order to explore two first-person approaches.

Certain forms of reflective self-consciousness, especially a methodological hyper-reflective or introspective consciousness, where attention is focused on one's own consciousness, are characteristic of abstract behaviour in detached states and pathologies, as well as in philosophical reflection and certain limit situations like fatigue and illness. In contrast, in most pragmatically and socially contextualized behaviour, when our attention is directed toward the external environment and we are involved in meaningful activity, both our reflexive and reflective consciousness are in many respects non-operative. Indeed as Csikszentmihalyi (1978) has shown, when people have been totally involved in an activity (e.g. rock climbing, problem solving) they retrospectively report that they were aware of the immediate situation but say that they cannot report the content of their conscious awareness at the time.²¹ Our intentionality is directed to things and other people; our consciousness is immersed in our projects. If we attempt to turn our reflective regard from our projects to the structure of consciousness, or the self, we alter our intentional structure, and the self who had been immersed in those projects is now abstracted from them. We end up with something akin to an artificially produced dissociated state: a self which is engaged in a project of reflecting, and a self which has been reflectively abstracted from the situation that had engaged it. Precisely in this disengagement, the reflected self ceases to be itself. Is there a way to capture the pre-reflective self, which, in any particular situation, is caught up in a unity of action?

It might seem that our only access to this pre-reflective self is either just the kind of abstract self-reflection that we have criticized, or the third-person observations that we have just set aside. If this were the case, then our position would not be unlike a Kantian faith in a transcendental entity which is distorted in every attempt to capture it. We admit that even these detached kinds of access, scientific and philosophical reflection and third-person observation, do provide some information about the pre-reflective self, but that the information is both abstract and incomplete, and in that sense distorted. We want to suggest two other kinds of access. The first one involves proprioceptive and ecological self-awareness. The second is a form of first-person contextualized access that we call 'embedded reflection'.

Ecological self-awareness

Within Gibsonian psychology one finds the concept of a non-observational access to what Neisser (1988) has termed an ecological aspect of self (see Butterworth, 1995; 1998; Marcel and Dobel, under review). This involves the idea that the information that I receive about the world includes, implicitly, information about my own self (specifically about egocentric perspective and spatial embodiment). To whatever

[21] The notion that we are not conscious of our situation and actions while driving a familiar route in a car is unwarranted. Rather, our reflexive or reflective consciousness is not focused on these things at the time of the activity, and we cannot, therefore, consciously recollect them. While both first-order conscious experience and second-order reflexive consciousness may be necessary for later episodic memory, the lack of formation of episodic memories, or the inability to remember episodically does not imply an absence of consciousness at the time of the original event.

extent this information is part of conscious experience, for example, in the form of proprioceptive awareness and awareness of egocentric self-location, it provides some sense of myself as an experiencing organism. My perception of the world is at the same time shot through with information about my own embodied position in that world. Ecological information (from both exteroceptors and interoceptors) about perceptual perspective, embodiment, and motor activity not only facilitates motor control, it provides a basis for distinguishing between self and non-self.²² Although much of the detail about bodily position and movement *vis-à-vis* the environment, detail which is absolutely essential for motor control and physical action, is not conscious, whatever *is* conscious does not present itself as detailed information about various parts of my body. Rather, it manifests itself as an integrated or global sense of where I am spatially in relation to the immediate environment and what, in any particular situation, I am capable of doing. In effect, ecological access provides a pre-reflective sense of the self as a spatial presence and a set of capabilities.

When my attention or conscious activity is directed toward the environment or toward some project, the content of proprioceptive awareness, in this Gibsonian sense, tells me, for example, whether I am moving or staying still, whether I am sitting or standing, whether I am reaching or grasping or pointing, whether I am speaking or maintaining silence, whether I am thinking or not.²³ Proprioceptive awareness thus provides an immediate experiential access to my pre-reflective, embodied self, even as I, as an agent, am not reflectively seeking myself, but am engaged in pragmatically and socially contextualized action. This is precisely what cannot be fully grasped in approaches that proceed reflectively within situations where the ecological sense is overtaken or dominated by reflective consciousness. One reason for this inability, so far unmentioned, is that in states of reflexive self-consciousness one's focal attention is by definition focused on oneself. Removing one's attention from the world and from one's goals destroys or makes perceptually recessive the self that is expressed in action. There is nothing mysterious about this. It is simply that an agentive self, as it removes itself from action to reflection, cannot perceive itself (as acting in the original action). David Rosenthal (1993) has noted a related point: pre-reflective phenomenal states can be *expressible* without being *reportable*. Only the content of second-order reflexive states of awareness can be reported (Marcel, 1993). This suggests that insofar as there is a contextualized self (a self contextualized in and by perception and action) it is something which can be expressed in action, emotion,

[22] This ecological sense is present from the very beginning of life. See Gallagher and Meltzoff (1996) and Butterworth (1995).

[23] James Gibson defines this ecological perspective in the following terms.

Awareness of the persisting and changing environment (perception) is concurrent with the persisting and changing self (proprioception in my extended use of the term). This includes the body and its parts and all its activities from locomotion to thought, without any distinction between the activities called 'mental' and those called 'physical' (1987, p. 418).

We also want to include in the concept of an ecological aspect of self what Neisser (1988) terms the interpersonal aspect of self. That is, we live in an environment that is not only physical but also social, and in social contexts we are constantly (but not necessarily consciously) provided with information about ourselves from the way others react to us. This represents a conflation of Neisser's two kinds of self which in other contexts are quite useful to keep apart. The idea that the self is unavoidably interpersonal is nicely captured in Varela's thought that the self 'cannot be separated from the distributed, multiple others which are our inescapable human ecology' (1996, p. 340, note 7).

or in certain attitudes, but not necessarily something which can be reported. To be reported it would need to be encountered in Hume's sense, that is, it would require second-order reflexive awareness. Reflexive awareness, however, necessarily involves an experiential separation between that of which one is aware (the object of reflexive consciousness) and the subject of awareness. This is the very activity which removes the self from the original action and decontextualizes it. The reflecting agent, which *expresses* itself in the action of reflecting, can only *report* on a self that is no longer in action.

The ecological self discussed above is a self that is perceptually specified. There is, however, another but related kind of embedded self-awareness, where the self is a more positive presence, but where the self is still not experienced as an objective 'Me' as it appears to a subjective 'I' (i.e., a self-awareness which is not a second-order reflexive awareness of a self separate from the perceiver). This is the sense of agency. When one is aware of one's actions at the time of acting, one experiences them as owned, as one's own. One does not experience them as unowned or as another's. In this respect one experiences oneself; and this is the case when we are involved in our actions, as opposed to being reflexively or retrospectively aware of them. The online sense of agency is thus complementary to the ecological sense of self in perception. Marcel (forthcoming) has argued that in one's immediate phenomenology during action, the owner or agent of action is not represented as separate from the action, but is an intrinsic property of action itself, experienced as a perspectival source. Findings in two areas of research concerned with awareness of bodily action and with felt ownership of bodily action point to a common conclusion.

First, experimental research on normal subjects by Marcel (ongoing) and by Jeannerod (Fourneret and Jeannerod, 1998) suggests that awareness of one's action is based not so much on actual feedback from movement itself or from peripheral effort associated with such movement, but more on that which precedes action and translates intention into movement — high-level motor commands, experienced as 'tryings'. Further, research by Haggard (Haggard and Eimer, in press; Haggard and Magno, in press) which correlates initial awareness of action with recordings of the Lateralised Readiness Potential and with Transcranial Magnetic Stimulation, strongly indicates that one's initial awareness of a spontaneous voluntary action is underlain by the motor commands relating to the effectors to be used. That is, although the content of experience is the action, its source is in fact what lies between intention and performance. It is important to note that the cerebral area centrally involved in this is the Supplementary Motor Area.

Second, the prime and most compelling instance of experienced dis-ownership of action is Anarchic Hand Syndrome (della Sala *et al.*, 1994). In this neurological syndrome one hand acts purposefully and wilfully against the person's conscious intentions and their efforts to suppress it. The action, but not the bodypart involved, is experienced as another's. Such patients do not have a delusion that someone else *is actually* doing it, i.e., they do not take the experience for reality, but it is a genuine 'as if' experience. It should be noted that, in contrast with this, in Tourette's Syndrome and Obsessive-Compulsive Disorder, and often in schizophrenia, the action itself is experienced as owned but the source of the action, an intention or command, is what is dis-owned. Indeed stimulation of the Central Thalamic Nucleus produces hand movements, and subjects have no idea why they did them, but the actions are not dis-

owned (Hécaen *et al.*, 1949). Thus awareness and ownership of intention is not crucial to ownership of action *per se*. While lesions to the Basal Ganglia–Cortical loops are often implicated in Anarchic Hand Syndrome, the most frequent and common lesion site is the Supplementary Motor Area — as noted above, the site of activity found to be associated with normal awareness of action.

On the basis of such evidence, Marcel (forthcoming) suggests that immediate awareness and experienced ownership of action go together and lie in the specific ‘trying’ underlain by pre-action specifications. It is logically necessary that such specifications for bodily action are in an egocentric frame of reference, since they specify movements in space from the body’s points of origin. Furthermore, for commensurability with spatial targets, bodily disposition has to be represented in egocentric coordinates. In the phenomenology of the normal person such action specifications yield a perspectival sense of the source of action, in which spatial points of origin and the spatiality and physics of intended action are specified (a counterpart in action to the egocentric perspectivalness of perception). This amounts to ownership of action in spatial terms and an immediate sense of agency. This sense of ownership is different from the sense of ownership for sensations, bodyparts or thoughts. It should also be noted that in this case the owner, the self in action, is not a substantive entity but a source. In Anarchic Hand Syndrome, the patient is (neurologically) deprived of the phenomenology of *this* source of action, and a different route to action is involved. The patient’s awareness of the movements may be by means of internal proprioception, but his awareness of the wilfulness of the hand’s actions is external, through haptic contact with it via other of his bodyparts and through his inability to suppress it. This proposal also accounts for several aspects of anosognosia for plegia, since patients without any loss of proprioception *per se* nonetheless experience movement in their paralysed limb as they had intended it (Marcel *et al.*, under review).

This perspectival sense of self as agent is normally set against and supported by a long-term background sense of agency (Marcel, forthcoming). This latter involves not only awareness that one has intentions but crucially the sense that one’s actions, or rather one’s tryings, are *reliably effective*, i.e., that spontaneous movements do come about in accordance with intentions that conform to the physics of the environment and the body. However, the main point is that the data suggest that self-awareness, experienced as ownership of action, is an intrinsic property of action itself, and is perspectival. It is thus complementary to the ecological sense of self in perception, and likewise involves no reflexive or detached self-consciousness.

Ecological self and temporal extension

Ecological self-awareness is normally considered to be momentary, providing a sense of posture or movement at any particular instant. But proprioceptive and ecological awareness also must include a sense of self over time, a sense of self as temporally extended. Even if our bodily position and embodied activities are constantly changing, and in that sense, impermanent or non-persisting, ecological self-awareness gives us more than just a snap-shot profile of our posture, location and action. Implicit in this kind of self-awareness is a sense of *what I have just been doing*, and, of equal importance, *what I can do*, and *what I am just prepared to do*, a sense of capability which goes beyond the momentary. This sense of capability implicitly involves a con-

tinuity from past experience, since my capabilities are to some extent created and constrained by my past experience as well as my present situation. This is not episodic memory, which can provide a sense of personal continuity; the kind of continuity at stake in ecological self-awareness is a continuity that involves past learning and that is implicit in motor capabilities such as riding a bicycle or swimming. This sense of capability also involves a projection of possible movements or actions, which are constrained and enabled by the present position of my body or my present embodied activities. In this sense, ecological self-awareness involves not a temporal knife-edge experience, but a changing ‘specious present’ that opens in the directions pertinent to the actor’s intentional activities. William James’ (1890) notion of the specious present has been recast by contemporary psychologists as ‘working memory’ (e.g. Baddeley, 1990; 1992). We concur insofar as working memory is conceived of as attentional involvement in current projects, but not insofar as it is conceived of as a subpersonal representational mechanism.

Two related questions can be raised here. First, is the self that is expressed and realized in contextualized action temporally extended? Second, is the self in action *experienced* as temporally extended further than the immediate past and future?

To address the first question, the self realized and expressed in action is temporally extended insofar as one’s actions are informed not just by one’s individual procedural learning, but by one’s past experience, by beliefs, lasting attitudes, moral positions, by one’s personal knowledge, concerns, and practical interests. One’s actions are often so informed by reflection on and recollection of episodic memories and autobiography. More to the point, however, and despite the contribution of reflection and recollection, one’s actions do seem to be informed by such things nonreflectively. This is partly what we will refer to as ‘character’ (see below). It is perfectly true that self-image and the avowals of character witnesses are no guarantee of what one will actually do or be capable of. Nonetheless there is a degree of consistency in an individual that is captured by the notion of disposition. Such a notion does not necessarily require either conceptual representation or reflective consciousness.

Regarding the second question, there are two respects in which temporal extension of self may be experienced in embedded action. The first of these is a primitive sense of time reflected by the ‘aspectual’ use of verbs, separable from tense *per se*. McCormack and Hoerl (in press) point to the fact that prior to children’s appropriate use of verb tense and prior to evidence for episodic memory proper, they are sensitive to temporal aspect, which marks not only completion versus continuation of action but also the difference between actions that are punctate versus extended (e.g. hit vs swim). Consider the differences between: I do it (now), I am doing it, I do it (habitually), I am about to do it, I complete it, I have completed (just done) it, I initiate it. Many of these temporal aspects of action are made available to us and are differentiated by our ecological awareness discussed above. But others, such as the generic ‘I do *x*’ (i.e. habitually), suggest that the very doing of an action brings into the momentary proprioceptive awareness of the actor the sense that he knows how to do *x* (the sense of capability discussed above) and either generic knowledge that he has done it before or even episodic memories of the action.

Episodic memory involves a second respect in which we may have an awareness of our own temporal extension in action. As opposed to deliberate attempts to recall episodic memories, preparation for and performance of action may bring to mind other

instances of performance of that action by oneself. Even if one does not have a sense of specific times or occasions of such episodic memories of doings, they nonetheless give one a sense of other instances of oneself as the identical first-person agent — as the perspectival source of action. Such awareness is not of long-term continuity, but of re-emergence or re-identification. While we do not doubt that people have a detached conception and belief in their long-term continuity, we are doubtful that people have an immediate *sense* or *awareness* of such permanence, other than when engaged in detached consideration of their personal narratives.

A somewhat more reflective awareness of temporal continuity is provided when, within activity or the planning of action, we deliberately recall episodic memories or consider how to behave. In the final section we suggest that this kind of reflection may also be ‘embedded’.

The Ethical Self

What we want to call ‘embedded reflection’ is not the same as the hyperreflective or introspective consciousness we identified in previous sections as a form of abstract, decontextualized behaviour. We may state the difference in this way. Embedded reflection is a first-person reflective consciousness that is embedded in a pragmatically or socially contextualized intentional attitude and the corresponding actions. It involves the type of activity that I engage in when someone asks me what I am doing or what I plan to do. In such reflection I do not take consciousness or ‘the self’ as a direct or introspective object of my reflection; I do not suddenly take on the role of a phenomenologist or theorist for the sake of answering the question. Rather I start to think matters through in terms of possible actions. I treat myself (I discover myself) as an agent. In such situations, my attention is directed not in a reflective inspection of consciousness as consciousness, but toward my own activities in the world where my intentions are already directed. Often my aim in such reflection is not to represent my ‘self’ to myself, as if it were a piece of furniture in my mind, but to continue certain actions or to explain myself in terms of my actions.²⁴

Korsgaard (1991) distinguishes between *being engaged in a conscious activity* and *being conscious of an activity*. If, for example, I move across the room in order to pick something up and hand it to another person, I am engaged in a conscious activity of voluntary movement and I know what I’m doing. If forced to express it I may say ‘I’m getting that book for my friend’. Part of that conscious activity may include an aspect of embedded reflection, and I may be thinking to myself, as I move, that I had better get this book to show my friend what I was talking about. Embedded reflection in this case is part of my engagement in the conscious activity of getting the book. If my actions call for a momentary or ongoing consideration of my intentions, this sort of reflective consciousness does not necessarily involve an interruption of or detachment from action. Certainly I may be engaged in getting the book and I may be simultaneously formulating a commentary (‘Listen, I’m just going to get that book in the

[24] Much of what we say in this section resonates well with points made about ‘embodied reflection’ by Varela *et al.* (1991). Invoking a Buddhist tradition, they suggest ‘a change in the nature of reflection from an abstract, disembodied activity to an embodied (mindful), open-ended reflection’ (p. 27; also see pp. 27–31). Like their work, our remarks here are also informed by the phenomenological tradition. The notion of embedded reflection echoes Heidegger’s and Merleau-Ponty’s critiques of the detached attitude.

next room') which may even serve to improve my actions, or clarify my contextualized intention. For example, not infrequently we may arrive at a destination, knowing that we came for some specific purpose, but unable to remember precisely what that purpose was. In some such cases, we may forget our intention precisely because we are so immersed in the immediate action. When, however, we do not lose track of our intention, embedded reflection can be reflexively directed to the intention itself. In such cases, embedded reflection can assist in keeping our intentions accessible, not as certain contents for epistemological investigation, but as pragmatic guides to our actions. Within this kind of self-reflection or reflection on my intentions, I would not necessarily be reflecting on the fact of my moving — on how my legs are moving or on how my arm reaches for the book — nor would I be thinking of the fact that my thoughts about the book are indeed conscious and perhaps organized in a successive way, etc. All of these things constitute the possible subject matter for a more abstract, theoretical or phenomenological reflection, but that would be something more than and different from embedded reflection.

We are not claiming that embedded reflection is such that on its own it could provide a theory of the self, as, for example, one might claim for a scientifically or phenomenologically trained reflection. Indeed, a theoretical model of the self is always going to be something more abstract than the kind of insight provided by embedded reflection. It is the nature of theory to involve generalization and to move beyond particulars. If reflection is embedded, it is embedded precisely in the particulars of an individual's life. We suggest, however, that theoretical models of the self ought to start closer to the level of the embedded reflective access we have to our own action and experience, rather than with detached introspective reflection on the nature of consciousness or self. Furthermore, an investigation of the self or personal identity that takes its bearing from embedded reflection and action is not necessarily naive, or uninformed by scientific or systematic considerations. One can educate embedded reflection and make it rigorous and systematic, even if it cannot be theoretical in the sense of being generalizable. Indeed, we find indications of this in long-established and sophisticated traditions of ethical thought.

One kind of embedded reflection is moral deliberation, in which I attempt to answer the question (posed by myself or by others): What do I plan to do? This kind of deliberative reflection includes what Charles Taylor (1985) calls 'strong evaluation,' a reflection on one's own desires and beliefs in which we evaluate our desires and not just the objects that we desire. Nonetheless, when we evaluate our desires, beliefs, and intentions in this way, we do not evaluate them as psychological or epistemological entities, nor do we report them as contents of our consciousness as such. Rather we evaluate them in their transparency, that is, in terms of our commitment to acting on them.²⁵ Taylor importantly points out that such strong evaluations are 'anchored' in feelings, emotions and aspirations, and can move us to take action about the way that we take action. Such reflective evaluation is embedded in what we inherently take to be meaningful and important, in our purposive designs for life, in what we intend to do. When in such evaluative explication we define or redefine our desires

[25] This notion of transparency derives from Wittgenstein and is discussed in Evans (1982) and Moran (1997). The idea is that, when I am asked what I believe about *x*, I respond by talking about *x*, not about the psychological or epistemological status of my belief. That is, we are aware of the content of our beliefs rather than of our beliefs *as beliefs*.

and intentions, we make constant reference to the way we are living or want to live our lives. It is at this level that episodic memory and certain emotions (especially emotions with a self-reflexive and temporal aspect, such as guilt, pride, hope or relief) can inform embedded reflection and provide a sense of continuity over time, and a sense of responsibility for actions.

To the extent that I come to realize that I am responsible for my actions, and feel so, I recognize a self that I have been, and that I continue to be, or that I want to transform by my actions. Insofar as I am willing to accept the consequences of my actions, or even insofar as I attempt to avoid the consequences of my actions, I am not only making reference to a pragmatically embedded self that is relatively integrated, i.e., dispositionally consistent over time, but I am able to provide an account of that self in the form of a reflective justification, a narrative, or an autobiography.²⁶ To be sure, such an account will be an explanation or interpretation, but it will be one that is tied to the reality of my own actions, and in that sense it will be a model of my own self. The referent of such reflection and expression is precisely a self that is instantiated in actions across a variety of personally and pragmatically contextualized situations. This notion of the self is best thought of as an *ethical or pragmatic identity*, for it is one that develops within a way of life (an *ethos*) and on a personal level.

Korsgaard suggests that ethical identity can be thought of as a unity of consciousness, but not in the sense of a unified collection of experiences. ‘The unity of consciousness consists in one’s ability to coordinate and integrate conscious activities’ (1991, p. 329). Perhaps we should not call this a unity of consciousness so much as a unity of active life. It is a unity in the particular way that we habitually tend to be engaged in *intentional actions*, rather than a unity of our consciousness of *such activities*. The unity here pertains to agency and capability for action, and involves, not a reflective retrospection that bestows identity on past experiences, but a prospective, embedded reflection that accompanies intentional action and helps to constitute capability. To be sure, deliberation among alternatives and the evaluation of their implications, happen, but not on an abstract theoretical level. Rather, deliberation goes on in the light of one’s desires, beliefs, ambitions, hopes, and attitudes toward other people.

Ethical terms such as *integrity, disposition, constitution* and *character* may signify more appropriately than epistemological or ontological concepts, such as the unity of consciousness or permanence over time, the important aspects of personal identity. Indeed, the central issue for us is not the question of unity, or the possibility of one or many selves. There may be no secure unity of self to be found in ethical behaviour, and we can certainly find examples of ethically inconsistent behaviours within the same individual. By character or disposition we do not mean an essence or thing; nor are these concepts captured by nineteenth-century conceptions of motive and cause that are conceived to be bound up in a nexus that lies beneath the subject and that putatively unifies the phenomenal discontinuity of someone’s existence, experience and actions. Character is formed by and continues to be formed by experience and by actions. What we are describing as the ethical self (involving character, disposition, and action capabilities) is generated in and through action.

Outside of pragmatically and socially contextualized actions, the self can only be pictured in an abstract snapshot, through a kind of reflection or experiment that robs it

[26] Again, we do not claim that there is any immediate sense of unity of self or of long-term temporal continuity. Rather, in feeling responsible for our actions such unity and permanence is *implied*.

of a rich *ethos*. The self, insofar as this signifies who the subject actually is, comes to exist and is articulated in the social and personal pragmatics and in action contexts that are characterized primarily in personal and social terms. Who I am gets expressed (and is constituted) in my actions towards others and my interactions with them. The personal self comes to exist as such only within such pragmatically and socially contextualized behavioural situations, that is, in action contexts that are characterized primarily in personal and social terms. The vocabulary of self is one that applies to my social interactions, my ethical projects, my meaningful intentions, my moral responsibilities, and to my bodily movements and linguistic competencies *insofar as they are taken up in such situations*.

Certainly what we have said about ecological access and embedded reflection is not a complete characterization of the ethical self. Each self has a certain ‘depth’ that can be traced out across an individual’s own history (including upbringing and individual life experiences), dispositional attitudes and concerns, and traced through the effects that culture and particular constraints imposed by various other factors (such as language, class, gender and race) have on the individual’s practical interests, projects and goals. To the extent that individual desires, as well as cultural practices and the performance of social roles, are developed and enacted most fully in behaviour that is defined by personal and social dimensions, the fuller picture would still require that the notion of self be described in vocabulary that is intrinsically related to pragmatically and socially contextualized behaviour. Who I am is most often tied to what I do, or what I can do, or how I express myself in such contexts, even if such actions and expressions cannot capture the full depth of the self.

One finds only incomplete or distorted aspects of the self in isolated or decontextualized movements, in physical mechanisms or functions that can be dissociated in an abstract fashion, or in metaphors such as the flow of consciousness or momentary Strawsonian ‘pearls’. Theories that start with abstract behavioural levels base their models of the self on partial, and sometimes distorted aspects of the self — aspects of body image, mechanical movements, syntax and grammar, the starts and stops of a conscious stream, but not on moving and speaking persons. In pragmatically and socially contextualized interactions, however, we find persons who are immersed in purposive projects and are conducting meaningful communication.

Our intention here has not been to work out the details concerning the structure of the ethical self, or how ecological aspects relate to the relatively integrated agent referenced in embedded reflection. We have not tried to explicate precisely how abstract reflective consciousness or psychological experimentation may provide positive snapshots of a self that is more, or less, integrated than its representations. Our intention, rather, has been to indicate an appropriate starting point for building a more adequate and less abstract model for explaining selves.

References

- Baddeley A.D. (1990), *Human Memory* (Hove, UK: Lawrence Erlbaum).
 Baddeley, A.D. (1992), ‘Working memory’, *Science*, **255**, pp. 556–9.
 Bisiach, E., Berti, A. and Vallar, G. (1985), ‘Analogical and logical disorders underlying unilateral neglect of space’, in *Attention and Performance XI*, ed. M.I. Posner and O.S.M. Marin (Hillsdale, NJ: Lawrence Erlbaum).
 Bisiach, E. and Geminiani, G. (1991), ‘Anosognosia related to hemiplegia and hemianopia’, in *Awareness of Deficit after Brain Injury*, ed. G.P. Prigatano and D.L. Schacter (Oxford: Oxford University Press).

- Brouchon, M., Joannette, Y. and Samson, M. (1986), 'From movement to gesture: "here" and "there" as determinants of visually guided pointing', in *Biological Foundations of Gesture*, ed. J.L. Nespoulos, A. Perron, and R.A. Recours (Hillsdale, NJ: Erlbaum).
- Butterworth, G. (1995), 'An ecological perspective on the origins of self', in *The Body and the Self*, ed. J. Bermúdez, A. Marcel & N. Eilan (Cambridge, MA: Bradford/MIT Press).
- Butterworth, G. (1998), 'A developmental-ecological perspective on Strawson's "The Self"', *Journal of Consciousness Studies*, **5** (2), pp. 132–40.
- Crick, F. (1994), *The Astonishing Hypothesis: The scientific search for the soul* (London: Simon & Schuster).
- Csikszentmihalyi, M. (1978), 'Attention and the holistic approach to behaviour', In *The Stream of Consciousness*, ed. K.S. Pope and J.L. Singer (New York: Plenum Press).
- Cumming, G.D. (1972), 'Visual perception and metacontrast at rapid input rates', Unpublished DPhil thesis, University of Oxford.
- Damasio, A. (1994), *Descartes' Error: Emotion Reason, and the Human Brain* (New York: Grosset/Putnam).
- della Sala, S., Marchetti, C. and Spinnler, H. (1994), 'The anarchic hand: a fronto-mesial sign', in *Handbook of Neuropsychology*, vol. 9, ed. F. Boller and J. Grafman (New York: Elsevier).
- Dennett, D. (1969), *Content and Consciousness* (London: Routledge and Kegan Paul).
- Dennett, D. (1991), *Consciousness Explained* (Boston, MA: Little, Brown and Company).
- Dennett, D. (1992), 'The self as a center of narrative gravity', in *Self and Consciousness: Multiple Perspectives*, ed. F. Kessel, P. Cole and D. Johnson (Hillsdale, NJ: Erlbaum).
- Donaldson, M. (1978), *Children's Minds* (Glasgow: Fontana).
- Evans, G. (1982), *The Varieties of Reference* (Oxford: Oxford University Press).
- Fournier, P. and Jeannerod, M. (1998), 'Limited conscious monitoring of motor performance in normal subjects', *Neuropsychologia*, **36**, pp. 1133–40.
- Gadamer, H-G. (1976), *Philosophical Hermeneutics*, trans. and ed. David E. Linge (Berkeley, CA: University of California Press).
- Gallagher, S. (1992), 'The theater of personal identity: From Hume to Derrida', *The Personalist Forum* **8** (1992), pp. 21–30.
- Gallagher, S. (1998), *The Inordinance of Time* (Evanston, IL: Northwestern University Press).
- Gallagher, S. and Meltzoff, A. (1996), 'The earliest sense of self and others: Merleau-Ponty and recent developmental studies', *Philosophical Psychology*, **9**, pp. 213–36.
- Gazzaniga, M.S. (1978), 'On dividing the self: Speculations from brain research', in *Neurology: Proceedings of the 11th World Congress of Neurology*, ed. W.A. deHartog Jager *et al.* (Amsterdam: Excerpta Medica).
- Gibson, J.J. (1987), 'A note on what exists at the ecological level of reality', in *Reasons for Realism: Selected Essays of James J. Gibson*, ed. E. Reed and R. Jones (Hillsdale, NJ: Erlbaum).
- Goldstein, K. (1940), *Human Nature in the Light of Psychopathology* (Cambridge, MA: Harvard University Press; reprinted by Shocken Books: 1963).
- Haggard, P. and Eimer, M. (in press), 'On the relation between brain potentials and the awareness of voluntary movements', *Experimental Brain Research*.
- Haggard, P. and Magno, E. (in press), 'Localising awareness of action with Transcranial Magnetic Stimulation', *Experimental Brain Research*.
- Hécaen, H., Talairach, J., David, M. and Dell, M.B. (1949), 'Coagulations limitées du thalamus dans les algies du syndrome thalamique: résultats thérapeutiques et physiologiques', *Revue de Neurologie*, **81**, pp. 917–31.
- Hume, D. (1739), *A Treatise of Human Nature*, ed. L.A. Selby-Bigge (Oxford: Clarendon Press, 1888, 1975).
- Husserl, E. (1991), *On the Phenomenology of the Consciousness of Internal Time (1893-1917)*, trans. J. Brough. Collected Works IV (Dordrecht: Kluwer Academic).
- James, W. (1890), *The Principles of Psychology* (New York: Dover Publications, 1950).
- Jeannerod, M. (1987), *Neurophysiological and Neuropsychological Aspects of Spatial Neglect* (Amsterdam: Elsevier).
- Kihlstrom, J.F. (1993), 'The psychological unconscious and the self', in *Experimental and Theoretical Studies of Consciousness*, Ciba Foundation Symposium # 174 (Chichester: John Wiley & Sons).
- Korsgaard, C.M. (1991), 'Personal identity and the unity of agency: A Kantian response to Parfit', in *Self and Identity: Contemporary Philosophical Issues*, ed. D. Kolak and R. Martin (New York: Macmillan).
- Leontiev, A.N. and Zaporozhets, A.V. (1960), *Recovery of Hand Function* (London: Pergamon).
- Marcel, A.J. (1983), 'Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes', *Cognitive Psychology*, **15**, pp. 238–300.
- Marcel, A.J. (1988), 'Phenomenal experience and functionalism', in *Consciousness in Contemporary Science*, ed. A.J. Marcel and E. Bisiach (Oxford: Clarendon Press).
- Marcel, A.J. (1992), 'The personal level in cognitive rehabilitation', in *Neuropsychological Rehabilitation*, ed. N. Von Steinbüchel, E. Pöppel and D. Cramon (Berlin: Springer).

- Marcel, A.J. (1993), 'Slippage in the unity of consciousness', in *Experimental and Theoretical Studies of Consciousness*, Ciba Foundation Symposium # 174 (Chichester: John Wiley & Sons).
- Marcel, A.J. (forthcoming), 'The sense of agency: Awareness and ownership of actions and intentions', in *Agency and Self-Awareness*, ed. J. Roessler and N. Eilan (Oxford: Oxford University Press).
- Marcel, A.J. and Dobel, C. (under review), 'Confusing frames of reference in real space and in imagination — the stabilizing effect of perceptual input'.
- Marcel, A.J., Tegnér, R. and Nimmo-Smith, I. (under review), 'Anosognosia for plegia: Specificity, extension and partiality of bodily unawareness'.
- McCormack, T. and Hoerl, C. (in press), 'Memory and temporal perspective: The role of temporal frameworks in memory development', *Developmental Review*.
- Merleau-Ponty, M. (1962), *Phenomenology of Perception*, trans. Colin Smith (London: Routledge and Kegan Paul).
- Moran, R. (1997), 'Self-knowledge: Discovery, resolution, and undoing', *European Journal of Philosophy*, 5 (2), pp. 141–61.
- Nagel, T. (1975), 'Brain bisection and the unity of consciousness', in *Personal Identity*, ed. John Perry (Berkeley, CA: University of California Press).
- Nagel, T. (1986), *The View from Nowhere* (Oxford: Oxford University Press).
- Neisser, U. (1988), 'Five kinds of self-knowledge', *Philosophical Psychology*, 1, pp. 35–59.
- O'Shaughnessy, B. (1980), *The Will*. 2 vols. (Cambridge: Cambridge University Press).
- Parfit, D. (1984), *Reasons and Persons* (Oxford: Oxford University Press).
- Poizner, H., Klima, E.S. and Bellugi, U. (1987), *What the Hands Reveal about the Brain* (Cambridge, MA: MIT Press).
- Quinton, A. (1975), 'The soul', in *Personal Identity*, ed. John Perry (Berkeley, CA: University of California Press).
- Radden, J. (1998), 'Pathologically divided minds, synchronic unity and models of the self', *Journal of Consciousness Studies*, 5 (5–6), pp. 658–72.
- Ramachandran, V.S. and Hirstein, W. (1997), 'Three laws of qualia: What neurology tells us about the biological functions of consciousness', *Journal of Consciousness Studies*, 4 (5–6), pp. 429–57.
- Rosenthal, D.M. (1993), 'Thinking that one thinks', in *Consciousness*, ed. M. Davies and G.W. Humphreys (Oxford: Blackwell).
- Rovane, C. (1998), *The Bounds of Agency* (Princeton: Princeton University Press).
- Sartre, J-P. (1957), *Transcendence of the Ego: An Existentialist Theory of Consciousness*, trans. Forrest Williams and Robert Kirkpatrick (New York: Farrar, Straus and Giroux).
- Sheets-Johnstone, M. (1999), 'Phenomenology and agency: Methodological and theoretical issues in Strawson's "The Self"', *Journal of Consciousness Studies*, 6 (4), pp. 48–69.
- Sperry, R.W. (1968a), 'Mental unity following surgical disconnections of the cerebral hemispheres', *The Harvey Lectures*, Series 62 (New York: Academic Press).
- Sperry, R.W. (1968b), 'Hemisphere disconnection and unity in conscious awareness', *American Psychologist*, 23, p. 723.
- Strawson, G. (1997), '“The Self”', *Journal of Consciousness Studies*, 4 (5–6), pp. 405–28.
- Taylor, C. (1985), *Human Agency and Language: Philosophical Papers I* (Cambridge: Cambridge University Press).
- Teasdale, J.D. (in press), 'Emotional processing, three modes of mind, and the prevention of relapse in depression', *Behaviour Research and Therapy*.
- Varela, F. (1996), 'Neurophenomenology: A methodological remedy for the hard problem', *Journal of Consciousness Studies*, 3 (4), pp. 330–49.
- Varela, F., Thompson, E. and Rosch, E. (1991), *The Embodied Mind* (Cambridge, MA: MIT Press).
- Weigl, E. (1961), 'The phenomenon of temporary deblocking in aphasia', *Zeitschrift für Phonetik, Sprachwiss und Kommunikationsforschung*, 14, pp. 337–64.
- Weigl, E. and Kreindler, A. (1960), 'Beiträge zur Auffassung gewisser aphasischer Störungen als Blockierungserscheinungen', *Archiv für Psychiatrie und Zeitschrift für Gesamte Neurologie*, 200, pp. 306–23.
- Wilkes, K.V. (1988), *Real People: Personal Identity without Thought Experiments* (Oxford: Clarendon Press).
- Wittgenstein, L. (1958), *The Blue and Brown Books* (Oxford: Basil Blackwell).
- Zahavi, D. and Parnas, J. (1998), 'Phenomenal consciousness and self-awareness: A phenomenological critique of representational theory', *Journal of Consciousness Studies*, 5 (5–6), pp. 687–705.