CHAPTER 16

EXTENDED COGNITION AND THE UNITY OF MIND

Michele Di Francesco

> Where then is the mind? Is it indeed "in the head", or has mind now spread itself, somewhat profligately, out into the world? […] Every thought is had by a brain. But the *flow* of thoughts and the adaptive success of reason are now seen to depend on repeated and crucial interaction with external resources. […] In a sense, then, human reasoners are truly distributed cognitive engines: we call on external resources to perform specific computational tasks, much as a networked computer may call on other net worked computers to perform specific jobs.[1]

1. TWO NOTIONS OF MIND

According to the model of the mind theorised by "post-classical" cognitive science, mental processes are embodied and distributed examples of cognitive processing.[2] Body and environment contribute to the achievement of our cognitive tasks in such a fluid and integrated way that they can be considered *bona fide* parts of cognitive agents. According to the extended model of cognition, the "mind" lies outside the body. What makes a piece of information cognitively relevant is the role it plays, and nothing prevents this role from being played by an external item. In turn, the extended model of cognition leads to an extended model of *subjectivity*, according to which the subject "is spread into the world".

The aim of this chapter is to compare the *extended mind* with the *personal mind*. The personal mind is the kind of mind we attribute to (human) persons, by means of folk psychological *intentional* language. The personal mind has two fundamental features: it is the locus of subjectivity and it is the locus of rationality. In other words, (i) it makes reference to a "subjective ontology" (Searle 1992), or, as I shall say, it designs a *subjective space* which requires intentional language in order to be described. Subjective space is characterised by the first person point of view; it expresses an *individual* perspective; it exhibits a peculiar unity and a phenomenology; and its contents are given to the subject in a privileged way. (ii) Personal mind is fundamental, too, when we explain human action; it characterises the "space of reasons",[3] with its normative and intentional features; it allows us to

understand rational actions as *rational*, and gives us the conceptual tools to characterise actions as *actions* (not merely happenings).

We should not identify the personal mind with the conscious mind. Unconscious mental phenomena may be part of personal mind. Neither should reference to personal mind imply a form of radical dualism. Personal mind may (should) be considered as the product of subpersonal processes whose nature is impersonal (third-personal) and causally explainable in terms of neural correlates, functional organisation, and so on (even if no explanation is available at present). In this case, I shall say that personal mind is *emergent* from subpersonal processes.[4]

In the following pages, I argue that mere causal-informational connections, which characterise a cognitive system in the extended mind paradigm, are not sufficient to explain the peculiar kind of unity which is essential to our notion of personhood and subjectivity. The unity of subjective space cannot be explained by means of the causal processes which constitute the extended mind. Those processes, in fact, are blind to the boundaries between inner and environmental processes. If we want a slogan: we may have extended cognition, but there are no extended subjects. And in order to be able to explain this fact, we need the concept of personal mind.

2. THE EXTENDED MIND MODEL OF COGNITION

"Where does the mind stop and the rest of the world begin?" This question opens Andy Clark and David Chalmers' seminal paper *The extended mind* (Clark and Chalmers 1998—C&C hereinafter), where they advocate a peculiar form of *externalism*: "an *active externalism*, based on the active role of the environment in driving cognitive processes" (C&C, 7). The conceptual background of active externalism is the model of the mind theorised by the new cognitive science. According to the new paradigm, (human) cognitive agents have a strong tendency "to lean heavily on environmental support"—making use of various instruments which vary from pen and paper to the nautical slide rule, "and the general paraphernalia of language, books, diagrams and culture" (C&C, 8). Environmental support allows us to perform *epistemic actions*, which are a kind of (cognitive) activity in which "a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognising as part of the world" (ibid.). According to C&C, this allows us to consider the world as part of the cognitive process.

One may think that, while brain-based cognitive resources are designed by nature to allow humans to obtain their cognitive goals, external resources are only contingently available (a subject may forget her pen, or her notebook). C&C reject this kind of objection, claiming that mere contingency does not rule out cognitive status. The reason is that from the extended cognition perspective, what matters is the causal coupling between the organism and its environmental resources, and external features are "just as causally relevant as typical internal features of the brain" (C&C, 9). As we shall see, this conclusion may be acceptable in case of extended *cognition*, but it faces serious problems when applied to the extended *mind*.

It is important to note that, despite the obvious differences, nothing prevents classical and embodied cognitive sciences from sharing some metaphysical assumptions about the nature of mental processes. Even if new cognitive science stops thinking of the mind/brain as a (self contained) computer, its main assumptions are still compatible with the fundamental idea according to which *information processing* is the basis of thought.[5] The difference from the old paradigm is that—this time—part of processing takes place in the environment, and/or takes advantage of the physical constitution of the involved cognitive system.

According to the extended model of cognition, fundamental aspects of the mind (such as beliefs and desires) lie outside the body:

> There is nothing sacred about skull and skin. What makes some information count as belief is the role it plays, and there is no reason why the relevant role can be played only from the inside of the body. (C&C, 14).

In this sense, the extended mind is a natural development of a conception of cognition based on information processing, which is not sensitive to the internal-external distinction. It adopts the notions of information processing and explanatory role as key concepts. This entails the following: (a) when we explain a cognitive activity, only causal processes are involved; (b) causal processes (in particular information processes) are blind to the boundaries between inner and environmental processes because mind, body and environment are part of an integrated system.

We have then to re-design the boundaries between cognitive systems and the world. Even in a more cautious analysis of the issue, one of the authors goes so far as to claim that such a new picture of the mind "threatens to reconfigure our fundamental self-image by broadening our view of persons to include, at times, aspects of the local environment" (Clark 1997, 214). In the following pages I would like to show that this reconfiguration is not necessary: persons are not at risk of spreading into the world. There is at least one important sense in which the coupling between persons and their extra-bodily cognitive resources is a relation between two different kinds of phenomena.

3. COGNITION, CONSCIOUSNESS, AND THE EXTENDED MIND

An obvious criticism of C&C's view is that external cognitive processing has no phenomenological content. There is nothing that "it is like to be" if you are an external support for cognition: why then call such states "mental"? They should be considered *tools* that a mind uses, and not part of the mind itself. The authors' reply is that the objection confuses the mental with the conscious. While it is implausible that our consciousness goes outside our body, there are many cognitive processes (connected with memory, language, skill acquisition—C&C, 10) which are not conscious. In fact, the distinction between mental and conscious is a commonplace in contemporary cognitive science, so any criticism of the extended mind model

based on the absence of consciousness in the external processing is the symptom of a confusion.

However, the issue is more complex. To characterize a form of cognitive processing which takes place in the external environment as "mental" suggests that it is in some way homogeneous with other kinds of processes which we normally call "mental". But this is disputable. There is no doubt that, if we adopt the extended mind paradigm, there is *causal homogeneity* among internal and external processes. But even the authors admit that the simple presence of a causal connection is not enough to secure the existence of a *mental* phenomenon. It would be silly to say that I know all the facts in the *Encyclopaedia Britannica* "just because I paid monthly installments and found space for it in my garage" (Clark 1997, 217). Note that the same problem arises when we compare mental activity and neural activity. There are many neural processes which are causally necessary for the existence of mental phenomena, but that are not called "mental": presence or absence of a certain neurotransmitter, the level of myelin on the axons, the glucose rate in the blood are obvious examples.

So, even if we confine ourselves to a strictly cognitive notion of mind and consciousness (ignoring for the moment any reference to the phenomenological level[6]), it is natural to distinguish between genuine mental processes and processes that are causally relevant for the emergence of the mind, but are not part of the mental realm themselves.

Note that we may well call "mental" many forms of cognitive processing that are not accessible to consciousness, but such processing should be strictly related to "conscious" processing. In fact, when we deal with the sub-personal level operating in the brain, the outputs of brain-realised cognitive processing become accessible to our conscious mind in a more intimate and direct way than those processed by external devices. And the nature of this intimate relation has a direct influence on the way our conscious mind works and interacts with other parts of our cognitive system. This point is obviously strengthened when we take into account subjective aspects of consciousness. Phenomenological content may affect subsequent mental processing;[7] consciously experienced beliefs are among the building blocks of "the space of reasons", and are thus essential to our explanation of human agency.

For brevity's sake, we may introduce the distinction between "subpersonal" and "nonpersonal" cognitive processing. Both are in a sense "external" with regard to the personal mind (generally speaking, subpersonal brain processes are not directly accessible to the subject), but they entertain a different relation to the personal mind. The crucial point is that, when the contents of subpersonal processing based on *internal* support have access to consciousness, they show different properties from those based on external supports. If I write a note in my notebook and later I read it, the note is presented to me in a different way from a memory. For example, they don't exhibit *immunity from error through misidentification*.[8] When I remember something I don't ask to myself: "here there are some memories; are they *my* memories?" (if I do, this may be a symptom of mental disorder). On the contrary, I may well ask myself "is that annotation mine?". Apart from schizophrenia *et similia*, I don't ask "are these thoughts in my mind *my*

thoughts?", while I may ask myself "is that handwriting *my* handwriting?". We shall return to this point later.

The "mineness"[9] of experience introduces the issue of the unity of the mind, the fact that conscious states are normally part of a unified conscious field *and* that "this total state is not *just* a conjunction of conscious states. It is also a conscious state in its own right".[10] *Mineness* and unity are also related to the *perspectivalness* of experience: the fact that our conscious experiences take the self as their common centre. Or, to use another metaphor, the fact that the subjective space where every conscious experience is located has relational structure, with the self as one of the *relata*—the intentional structure of conscious experience takes the self as a stable element, relating with changing object-components.[11]

The ways in which mental contents are available to the *personal* mind entail an immediate matching between the subject and her "internal" states. That is, when the subject has (conscious) access to the contents of her personal mind, she immediately and directly experiences them as part of *her* experience. When she entertains I-thoughts[12] (thoughts that are typically expressed by the first person pronoun "I"), and she thinks to herself in the first person perspective, there is no room for mistakes of identification. For the subject to have an experience is to immediately and directly experience it as *her* experience.[13]

Of course we may use the notes in our notebooks to remember something, simply by reading them. When we acquire information by conscious perception we may take the result of external processing as input for our personal mind. But the connection between the information so acquired and the *personal* mind is quite different from the connection which originates from subpersonal processing. In the latter case (a) the content is immediately and directly available to the subject; (b) it exhibits a co-presence relation with the other (conscious) mental states of the subject; (c) it may be taken as a non-inferential starting point for a judgement.[14] All these features of personal mind are not present in the extended mind.

Right now, it is easy to imagine a criticism to my line of thought: instead of rejecting the extended model paradigm, we have rather changed the subject, switching from extended cognition to consciousness. I would like to stress that my point is *not* that extended cognition is impossible. The purported irreducible subjective character of the (conscious) personal mind is not *direct* evidence against the extended mind model, since that model deals with a more general notion of cognitive processing. My (first) point is that there is a *gap* between extended mind and personal mind, and that the conceptual resources available to the extended mind model are not sufficient to fill the gap. And (my second point) if we cannot fill the gap, there is no future for the notion of the extended subject.

4. FROM EXTENDED COGNITION TO EXTENDED MIND?

The existence of a gap between extended and personal mind can be argued to make clear that there are two *kinds* of informational processing connected to our cognitive capacities. The first is constituted by those *subpersonal* processes that offer direct input to the personal mind. The second is the wider set of computational

activities that are causally relevant to the achievement of a determinate cognitive goal. This wider set of cognitive processing may make use of external support, and in this case interacts only indirectly with the personal mind.

We should note that the interesting distinction here is not internal vs. external, or biological vs. mechanical. Let us suppose that it were possible to install in our brain electronic implants which made a certain kind of information accessible to our conscious mind, in the immediate and direct way that *biological* subpersonal processes do.[15] In this case I think that we should consider them as examples of the first kind of information processing, since they would offer direct input to the personal mind.[16]

There may be many theories of the relation between subpersonal processing and the personal mind. A moderate version tries to limit the "dualistic" consequences of the personal/subpersonal distinction. For example it may present it as a "change of destination" of the same input, as in the following model, proposed by Gareth Evans:

> We arrive at conscious perceptual experience when sensory input is not only connected to behavioural dispositions […]—perhaps in some philogenetically more ancient part of the brain—but also serves as the input to a *thinking, concept applying, and reasoning system*; so that the subject's thoughts, plans, and deliberations are also systematically dependent on the informational properties of the input. When there is this further link, we can say that the person, rather than just some part of his brain, receives and possess the information. (1982, 158).

Here, even if the personal level is a *new* emergent aspect, a systematic connection is required between the content of the thoughts exercised by the subject and the informational properties of the inputs.[17] More radical versions may choose to underline the distance between "the space of causes" and "the space of reasons".[18] The moderate and the radical approaches will offer us different treatments of the connection between thoughts and behaviour, but I am not interested in this aspect of the issue.[19] What is in question now is the emergence of the personal level from the subpersonal one.

In particular I claim that, while it is conceivable that the personal mind may emerge from subpersonal processes (from the "subpersonal mind"), it is very difficult to find an appropriate relation connecting extended and personal mind. The differences between the two ways of treating mentality are so great that we may suppose a kind of *incommensurability* between them.

An obvious reply could be that it must be possible to fix within the extended mind a proper subset of processes which can be considered as the emergence-basis of the personal mind. But the point is that it is not possible to *individuate* this subset by means of the conceptual resources of the extended mind paradigm. For the very reasons given by its proponents, the extended mind model invalidates and nullifies the internal/external distinction.[20] It is by *starting from the personal mind* that we

can discriminate the extended mind processes which are suitable candidates for the subset.

To clarify this point, let us consider the foundations of the unity of the extended mind. We know that it is based on causal-informational relations (those relations that are involved in the cognitive processing developed by the system). Is *any* relation good enough? Suppose that an apple hits my head, causing a thought about the law of gravity. Is the apple part of the process of thinking? It seems obvious that a positive answer raises problems—a fact that C&C acknowledge, as is apparent from their discussion of what we may call the "portability objection". The brain—the objection says—comprises "a package of basic, portable, cognitive resources" which may incorporate bodily actions, but which does not encompass "the more contingent aspects of our external environment" (C&C, 10-11). The authors' reply is that "mere contingence of coupling does not rule out cognitive status" (ibid.) for two reasons. The first is that in the future we may be able to plug technological modules into our brain (for example, a chip that increases our short-term memory). Our reply to this is that *if* such an implant produced a genuine form of memory which offered direct input to the personal mind, it would simply represent a non-biological development of the biological bases of subpersonal processes—and the result of its processing would not be external to the personal mind.

Quite a different situation occurs when we consider truly external devices such as pocket calculators, notebooks, etc. In this case, to answer the portability objection, C&C acknowledge that *reliable coupling* is required. To quote a related work by Clark:

> Mind cannot usefully be extended willy-nilly into the world. […] The notebook is always there—it is not locked in the garage, or rarely consulted. The information it contains is easy to access and use. The information is automatically endorsed—not subject to critical scrutiny, unlike the musings of a companion on a bus. Finally, the information was originally gathered and endorsed by the current user (unlike the entries in the encyclopaedia). (1997, 217).

These requirements sound reasonable. In fact, from my point of view, they try exactly *to mirror* at the causal level the phenomenological properties of the personal mind we have described above. We may describe their function by saying that they allow the passage from *cognitive processing* to *mind*. But from the extended mind perspective, I claim that they are completely *ad hoc*. Let us take the third requirement as an example: "the information is automatically endorsed – not subject to critical scrutiny". This is exactly what happens when the result of subpersonal processing becomes accessible to the conscious-personal mind of a subject. The resulting mental item is "just there", as a part of the integrated mind of the subject; poised for thinking, reasoning (and feeling). But why should such an uncritical attitude be extended to the content of the notebook? In fact, if I read a note I can doubt it, or I can reject it as unbelievable; I can suspend my judgement about

its paternity, and so on. I can carry out all these acts regardless of any explicit promise I may have made previously such as "I promise I'll never distrust my notebook entries".

In the extended mind model, the mind is a Parfittian entity,[21] a kind of club, collecting any processing that leads to the accomplishment of a cognitive task. And it is a democratic club: no processes can be excluded, according to the club rules. The restrictions described above mean to facilitate the passage from processing to mind, but they implicitly use a view of mind which comes from the "personal" side—which refers to our subjective inner life as first personally experienced. Of course "reliable coupling" is a useful biological requirement, but it has nothing to do with a process being an example of information processing, or not. The extended mind model needs a sub-set of processes which are constituents of a reasonably integrated *subject* (no apples within it). To strengthen their position, the extended mind's supporters have to navigate the dangerous waters of subjectivity. Is an extended subject possible?

5. FROM EXTENDED MINDS TO EXTENDED SUBJECTS?

In our target paper, C&C explicitly claim that they are not simply speaking of external processing: they want to describe extended *minds.* As we have seen at the beginning of section3, they make a distinction between experiences (which may be determined internally by the brain—C&C, 12), and other mental states such as belief, which can be constituted partly by the external environment.

We have already noted that it is not obvious that such a distinction is tenable, since the phenomenal component of our mental content may be relevant to the way we compute them. But in any case, beliefs are sufficient to make room for subjective spaces and rationality in our framework—shifting, so to speak, from cognitive agents to subjects. In human beings at least, beliefs can be consciously entertained, consciously endorsed or rejected, and they are referred to in our explanation of action. In other words, beliefs (and other propositional attitudes) have access to "the space of reasons"—and to its normative and intentional features, which are coherent parts of the personal mind's territory.

In any case, something is a belief for the extended mind paradigm if it conveys, in a reliable way, content that is causally relevant to the achievement of a cognitive goal. The reliability condition must be added to avoid purely contingent causal processes becoming part of the mind (as in the apple example above). So described, however, a belief is hardly definable as "mental", for the reasons I have proposed above. To characterize a vehicle of content as "mental" should require that it can be taken as input for the personal mind of a subject—exhibiting in this way the right connections with his phenomenology, rational agency and capacity for thought. The point is not terminological, but substantial: of course, we may call anything we want "mental". But, following the extended mind usage, we miss the distinction between subpersonal and nonpersonal cognitive processing. And the introduction of the abovementioned requirements to secure "reliable coupling" does not help.

To clarify this point we may look at the story of Otto and Inga (C&C, 12). Otto suffers from Alzheimer's disease, and carries with him a notebook which plays the role of his biological memory. Whatever new information Otto learns, it is written in the notebook. Let's suppose that he is told about an exhibition at the Museum of Modern Art, and that he decides to go to see it. Then he looks in the notebook for the Museum address, finds it, and so goes to the Museum. Now, if we compare Otto's story with that of Inga—a friend of his who heard about the exhibition and simply remembered the Museum address—we may think that Otto's notebook plays exactly the same cognitive role as Inga's biological memory. Especially if we introduce suitable restrictions to secure "reliable coupling": Otto constantly uses the notebook, he easily finds any information it contains,[22] he automatically endorses it,[23] and so on. "Otto's and Inga's cases seem to be on a par: the essential causal dynamics of the two cases mirror each other precisely" (C&C, 13).

Causal dynamics is not all, however. It fits pretty well with an enlarged version of cognitive processing, but it simply misses the distinction between the property of being an external item accessible to perception and the property of being an internal content accessible to thought. Both properties are causally relevant to explain Otto's behaviour; but they are, nevertheless, different in kind. Note that the point here is not the simple lack of a phenomenology associated with the retrieval of information in Otto (C&C, 16). The point is that the absence of a phenomenological dimension associated with the notebook entries offers evidence for the fact that these entries are not potential entries of Otto's personal mind—they do not contribute to making him the particular subject he is. The connection between Otto's mental content and the notebook entries is causal, but not *motivational*: it does not explain Otto's acts as *actions*. In other words, to be considered a *reason*, the notebook content should be assimilated within Otto's personal mind. And this is not the case. If Otto reads in the notebook "I want to go to the exhibition" and he doesn't go, this is *not* a case of weakness of the will. If the notebook contains incompatible projects, this does not amount to an interior conflict. If Otto reads "I can fly like Superman", but he doesn't believe it, he is not contradicting himself; and so on.

Of course, according to the requirements proposed to secure "reliable coupling", in many respects Otto *functions* as if the notebook were part of his personal mind. As we have seen, the reason is that the requirements mirror, at the empirical level, the *conceptual* requirements which define the personal mind (as it is shown by phenomenological investigation and logical analysis). But are constancy, accessibility and automatic endorsement which are acquired ad-hoc sufficient to say that we may have subjects who are "spread into the world"? According to the notion of subjectivity we adopted before, the answer is negative: Otto is not an extended subject who carries on fairly well thanks to his extra-mind prosthesis. He is an ill person, who can attain his cognitive goals by means of external support.

But perhaps this kind of criticism simply begs the question. Perhaps we should enlarge our notion of subjectivity to grant that integrated systems may be called "subjects". I don't think that this is a good idea: rationality, normativity and phenomenology are constitutive aspects of subjectivity which do not extend themselves outside personal mind. If we want to renounce them, we had better stop

speaking of *subjects*.[24] Let us suppose however that we may have extended subjectivity: "a coupling of biological organism and external resources" (C&C, 18); and, to avoid confusion, in our example, let's call this enlarged entity Super-Otto. Super-Otto is an enlarged (expanded) system which may be compared with Otto. Now, to demonstrate that they exemplify two different kinds of entities, I shall consider the very different ways in which they obtain their "mental" unity.

In fact, many aspects of the unity issue have been anticipated. When Otto exercises his perceptual abilities to obtain information from his notebook, the contents he acquires have a different connection from those contents available to him (the subject) as the emergent result of some subpersonal processing. In this latter case, they are part of a phenomenological and unitary mental space; they entertain conceptual, normative and motivational connections with other mental states; (when occurring as component of I-thoughts) they exhibit "immunity from error through misidentification"; they cannot be expressed by a definite description. The ways in which mental contents are available to the personal mind entail an immediate matching between the subject and his "internal" states. That is, we are entitled to speak of personal mind when having an (egological) experience is for the subject to experience it immediately and directly as *his* experience.[25]

None of this can be said about Super-Otto: all of the information he/it gains has identical status. He/it is a real Parfittian entity, which can incorporate any form of (reliable) information processing[26]—the only necessary glue being a causal connection. Of course we may try to reconstruct within his/its enlarged mental space the distinction between subpersonal and nonpersonal processing—thus explaining the difference between processes that affect the unity of "pre-mental" level from those which affect mere cognitive processing. But, again, this distinction is made starting from the personal level and looking "down" at the "subvenient" levels (which can be taken as the emergence bases of mental properties individuated at the personal level[27]). This means that the concept of personal mind is our only access to the distinction between (genuine) subpersonal processes which can be taken as the emergence-basis of a subjective space and the other nonpersonal informational processes which characterise the extended view of cognition. The extended mind paradigm is not enough to capture fundamental aspects of phenomena we call "mental".

6. CONCLUSION: OTTO, SUPER-OTTO AND THE PERSONAL MIND

If the reflections proposed here are convincing, the reader will be prepared to accept the idea that the perspective of personal mind is irreducible to the perspective of extended mind. Cognitive processing is not sensitive to the internal-external distinction and it is blind to the boundaries between inner and environmental processes. However, when we take mind, body and environment as part of an integrated system we lose sight of the *subject* and its *mental space*. But since personal mind phenomena do exist, the world-description we arrive at by means of the extended mind paradigm is incomplete.

Now a new question arises: let us concede that the perspective of personal mind is irreducible; but what if it is also *irrelevant*? In other words, are personal mind phenomena important data if we want to explore the "cognitive dynamics" generated by the relation of a (human) organism and its environment? Or are they just cognitive fossils, memories of a primitive stage of human development, of very limited importance when we try to understand the essential features that characterise our species today?

We may perhaps credit Clark (1997, 216) with a similar conception. Here Clark addresses the question as to whether the "putative spread of mental and cognitive processes into the world implies some correlative (and surely unsettling) leakage of the self into the local [[28]] surroundings?". Clark's answer is "Yes and No", and the "No" component depends on the fact that he credits the subject with conscious contents which supervene on the individual brain. But then he specifies that such conscious contents are not very important: they are "at best snapshots of the self considered as an evolving psychological profile" (ibid.). Being limited to just these conscious contents, we would miss the evolution of reason and thought, characterised by the intimate and complex interplay between mind, body and environment described by the extended mind model of cognition.

I think that this view contains more than a grain of truth when it stresses the limits of self consciousness, but that it seriously underestimates the importance and explicative richness of the personal mind. It also underestimates the personal mind's extension, which is not limited to ongoing conscious phenomena, but contains all the phenomena which can become accessible to the mind or (more generally) can serve as input to a *thinking, concept applying, and reasoning system.* From this perspective our personal mind is relevant in the explanation of behaviour not only because of the importance of emotive and sensitive factors in the construction of self and rationality.[29] Another reason is that the personal mind may internalise, so to speak, aspects of the cognitive environment—first of all language.[30] Cultural development in humans allows the personal mind to incorporate part of their social environment simply by concept- and language-acquisition. More literally, internalised language is part of personal mind, and internalised language makes our subjective world richer and deeper. The mental life of a language-speaker is qualitatively different. To give one example, let us suppose that we make a mistake that greatly affects our life. We take the wrong decision about a job, and, as a consequence, our personal life faces serious troubles. When we realise the consequences of our mistake, we may describe our mind as full of sorrow, grief, misery, remorse, sadness, regret, contrition, and so on. The more conceptual resources we have, the more fully we will be able to analyse our inner states. And the more fine-grained this analysis will be. And this is just the beginning. These mental states will motivate actions and conduct based on beliefs, opinions, desires and hopes, whose coherence and desirability we can evaluate in advance. In this sense, the mental life of a subject which entertains these kinds of mental states is not a static blueprint, but rather a dynamic reality. All these states of mind are detected by introspective ability powered or even made possible by language.[31] Note however that—differently from the extended mind model—here language is not an external

part of a boundary-less entity, but has been internalised to act as a proper part of an integrated mind.

I think this is a very important difference, and probably it is not by chance that Clark himself (1997, 198) wants to distinguish his position from Daniel Dennett's (1991) apparently similar claims. Dennett (1991, 219) says that our cognitive abilities depend only partially on our brain "hardware"—innate structure. More important are the ways in which culture and language affect the plasticity of the brain. Public language acts on a programmable brain, and so modifies its neural organisation. So—in Clark's words—for Dennett, public language is both a tool and a source of brain (mental?/cognitive?) reorganisation. In opposition to this, Clark proposes the extended mind paradigm, which is "inclined to see [language] as in essence just a tool—an external resource […]" (Clark 1997, 198).

On the contrary, I think that Dennett's model is useful to order to see how the personal mind is more than mere consciousness—and that it is *relevant*. Powered by internalised language (and other external practices, such as mathematics) personal mind has become rather efficacious and complex—very different from the extension-less locus of immediate consciousness depicted by the extended mind model.

Note that, if I agree with Dennett that our personal mind is *relevant*, I disagree with him in considering it *irreducible*.[32] As is well known, the self for Dennett (1991) is a virtual entity whose unity and coherence is fictitious, and which is reducible to subpersonal distributed and parallel processes. I think instead that we should think of the subject as an emergent entity, endowed with causal powers and the capacity to interact with the real world. The thesis of reducibility of the personal mind to subpersonal brain activity, however, is quite a different issue from the reducibility of the personal mind to the extended mind. To repeat it for the last time, if we adopt the extended mind stance, any distinction between mind and world vanishes. We cannot isolate—by means of a purely impersonal causal language—those processes that we take as internal to the subject from those that lie outside the realm of subjectivity (with its normative and intentional character). The very existence of a unitary subject is revealed only when we adopt the personal mind stance. To express it in a single sentence: causes alone cannot delimit the subject from the inside; but if we select the *relevant* causal chain by starting from the self then we are already outside causal language.

In saying this, I am more interested in the defence of the necessity of the personal mind than in the criticism of the idea of extended cognition itself. Contrary to other criticisms (which are for example discussed in Clark 2005), I consider extended cognition talk useful and insightful in many respects, especially when it casts light on the importance of social and cultural factors in shaping human cognition (and brain function). There is no reason to deny that extended *cognition* is a real and very important phenomenon. First, our cognitive tools change the way our mental tasks are developed and articulated dramatically. Second, socially powered mental phenomena

expand the capacities, aims and performance of our personal minds. As a consequence of their (biologically based) openness, human minds (and in general human beings) are open to technological complements. But to be so incorporated, cognitive technologies must produce their internal counterpart (external support is not enough). In this sense, even if external cognitive devices shape our personal minds, and our personal minds are part of the world, we—the subjects—are not spread into the world. So far as agency and thought are concerned, we are unified entities, individuals with a perspective on things—who are not only in the world, but also *have* a world; we are subjects of acts, and rational entities.[33]

NOTES

[1] Clark (1997, 68-69).

[2] See this volume, chapter 1, subsection 2.2.

[3] Sellars (1956); McDowell (1992).

[4] Emergentism claims that reality is structured in different levels and, when a given level reaches a certain degree of complexity, *new* properties emerge. These properties are not explainable and are not predictable on the basis of knowledge of their lower-level bases. They also bring into existence new causal powers, which can exercise their action upon lower levels. In fact there are different kinds of emergentism, but I am not taking any position about the *kind* of emergentism (radical or moderate) which should be adopted here. See Di Francesco (2005) for a discussion of various kinds of emergentism.

[5] Di Francesco (2002, 2004); Marconi (2001).

[6] We may adopt Block's (1995) distinction between access and phenomenal consciousness, and confine ourselves to access consciousness.

[7] This claim could be criticised by a defender of the epiphenomenal character of phenomenal consciousness. But I think that there is evidence of a causal role played by phenomenal properties. To propose just one example, it may be argued that phenomenal properties play a causal role in informing the higher order cognitive system about the nature and origin of the acquired information ("source monitoring"—see Johnson 1988; Johnson, Hashtroudi and Lindsay 1993). In this sense, they are *cognitively* relevant, since they may have causal efficacy in the process of evaluation of belief and decision making (see Zalla 1996, and this volume, chapter 14, section 5).

[8] Evans (1982, chapter 7).

[9] Metzinger (2004).

[10] See Bayne and Chalmers (2003, 27).

[11] I take the unity of the personal mind as a (phenomenological) fact which should be taken into account in our theorising about the mind. However, I admit that my analysis is rather sketchy and *prima facie* open to the criticism that new results in cognitive science and brain sciences may compel us to give up our phenomenological intuitions. Bayne (this volume, chapter 15) presents a detailed discussion of the many ways in which we may interpret the unity of consciousness thesis. He also shows that experimental or clinical case studies, or the adoption of certain general models of consciousness, may raise doubt about (various versions of) the unity thesis. The issue is rather complex, but I think there is room to defend the unity thesis if we adopt an emergentist picture of the mind, according to which cognitive disunion at the sub-personal level is compatible with (emergent) unity at the higher personal level. Personal mind unity (which is different from the unity of consciousness) is achieved by unconscious agencies; but it manifests itself as a real (and causally efficacious) phenomena—operating at a different level from that of its sub-personal basis.

[12] Evans (1982, chapter 7).

[13] Here again I take I-thoughts *as emergent* upon a set of abilities, information links among our bodily states, perceptual states, memory states and so on, whose complexity and divisibility into parts are not relevant to our acknowledgement of the presence of a unitary subject of experience. Of course, preconditions for the emergence of I-thoughts must be satisfied in order for I-thoughts (and subjectivity) to exist. But at *their* level of emergence, I-

thoughts and subjectivity exhibit properties which must be considered primitive and are constitutive of our idea of a subject of experience.

[14] This means that it is part of the space of reasons. Note that this typically holds for beliefs and other cognitive states, without taking into account phenomenal states as such.

[15] See Clark 2003, chapter 1, for a presentation of this issue.

[16] So we agree that "there is nothing sacred about skull and skin". The internal/external distinction is relative to the (personal) mind, not relative to the body.

[17] Evans (1982, 159).

[18] Sellars (1956); McDowell (1992); Brandom (2000).

[19] I also explicitly avoid any reference to scientific theories of consciousness – confining myself to the philosophical field. The reason is that there are not widely accepted scientific theories of consciousness. As concerns the two philosophical perspectives here described, it is perhaps of some interest that we may read the moderate and the radical versions as based on two different interpretations of the emergence relation. See Di Francesco (2005), for these interpretations.

[20] See however note 12.

[21] See Parfit (1984).

[22] A very difficult requirement to fulfil, and not only for an Alzheimer's patient—especially if we consider the huge amount of information the notebook would have to contain to be used as a viable alternative to our biological memory. See Marconi (2005) on this point.

[23] Here again we may doubt whether it is really possible to suspend the critical evaluation of the acquired information, when it comes by way of perception and reading, rather than mere presentation as content of the personal mind.

[24] Nothing that I say in this paper, in fact, shows that such an eliminativistic attitude is mistaken. See Metzinger (2003) for a very impressive attempt in this direction. But in this case I think we should admit that there are no subjects (and stop speaking of extended subjects; subjects spread into the world; and so on).

[25] Of course, Otto is an ill person. So sometimes he cannot achieve the sufficient integration of cognitive resources to act (and think) properly. And he can make use of technological aids. But in this case we say that his personal mind is impaired, and that he uses extra mental devices—not that his mind encompasses those devices, which are more like wheelchairs than transplanted legs.

[26] Marconi (2005) discusses an interesting mental experiment where a father's mind is part of the extended mind of a lazy daughter (who uses it as an automatic translator from Latin to Italian).

[27] In this sense, these subpersonal properties exhibit a form of ontological dependence on the emergent mental properties. This is why we can consider them, in a sense, mental. (See Di Francesco 2005, for the ontological analysis of this issue.)

[28] Why *local*? In the extended mind paradigm nothing prevents the self from being spread to the most remote regions of the world.

[29] The connection between reason and emotion is a widely accepted thesis in contemporary cognitive science and neuroscience. To drop one name, we may simply refer to Antonio Damasio's work in this field. See Damasio (1995, 1999, 2003).

[30] Differently from the extended mind causal model, this process is based on brain modifications connected to language acquisition and verbal competence, and it conforms pretty well to the idea of brain processes as the emergence-basis of thought and subjectivity.

[31] I shall not address here the question as to whether mental states themselves are created by language, even though, generally speaking, the relation between language and consciousness is probably quite relevant to our present issue.

[32] A nice set of arguments against Self-elimination can be found in Kennedy and Graham, this volume, chapter 17 (even though they are cautious with regard to Dennett's real position).

[33] I would like to thank Tim Bayne, Diego Marconi and Massimo Marraffa for comments on earlier versions of this paper.