THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Keeping the collectivity in mind?

Citation for published version:

Collins, H, Clark, A & Shrager, J 2008, 'Keeping the collectivity in mind?', *Phenomenology and the Cognitive Sciences*, vol. 7, no. 3, pp. 353-374. <https://doi.org/10.1007/s11097-008-9100-8>

Digital Object Identifier (DOI):

[10.1007/s11097-008-9100-8](https://doi.org/10.1007/s11097-008-9100-8)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Phenomenology and the Cognitive Sciences

Publisher Rights Statement:

© Collins, H., Clark, A., & Shrager, J. (2008). Keeping the collectivity in mind?. Phenomenology and the Cognitive Sciences, 7(3), 353-374. 10.1007/s11097-008-9100-8

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Keeping the Collectivity in Mind?

Harry Collins, Andy Clark and Jeff Shrager

Abstract

Collins accepts Clark's idea of the extended mind and says that the degree of flexibility accorded to the brain and mind on this model supports the notion of interactional expertise and the importance of language. Our relationship to the material scaffolding is, however, asymmetrical, so that the self is not just the sum of the parts. He claims that the idea that the individual is merely part of the collectivity or minds has long been known to sociologists of science so, if the extended mind comes as a surprise, cognitive psychology needs to change. Clark accepts the asymmetry point but argues that our metaphysics is in need of repair. Shrager responds that cognitive psychologists have long accepted that individuals are not isolated and agrees that the problem of socialisation is crucial. Collins argues that a reasonable metaphysics will have to keep a special place for humans and for the human collectivity. Clark responds that an asymmetrical ensemble is sufficient to explain the self. Collins responds that we know too little about the material world to pursue such a model to the exclusion of other ideas and that language has agency. We must 'keep the collectivity in mind,' as it were!¹

Keywords: Cyborg, extended mind, interactional expertise, collectivity, language, cognitive psychology, self.

¹ Though what follows is a robust exchange of views it is also a cooperative effort, authors communicating 'backstage' with each other to try to make the disagreements as clear and to the point as possible.

The Cruel Cognitive Psychology (by Harry Collins)

***The Cruel Sea*²**

When I was a kid in the 1950s I was much moved by 'The Cruel Sea,' a film about Britain's war against the U-boats. Shot in black and white, it is about failure, frustration and futile waste of life. A couple of years ago I watched it again and suddenly realized that more than about war it was about social class. The newest recruit to the wardroom of the 'Compass Rose' immediately signifies his working class origins with loud and boorish mannerisms and enthusiasm for sausages, which he calls 'snorkers.' The rest of the officers speak quietly and disdain sausages. When times get tough it is, of course, the loud snorker-lover who cracks while the quiet snorker-disdainers do their duty with stiff upper lips. Why didn't I notice this when I was a kid? I was focused on the echo-sounder, the depth charges, and the death of the innocent; I had neither the peripheral vision nor the right perspective for class analysis.

I believe that in five, ten, or, the way things are going, a hundred years, we are going to look back at Andy Clark's (2003) book and say 'why didn't people see what it really meant?' The book purports to be about our cognitive relationship with the physical world but I am going to suggest that its true meaning lies at its periphery. It is its silences that speak about the curious state of modern cognitive psychology. In the second part of these remarks I will try to bring nearer the day when we grasp this true meaning.

Let me start, however, by saying that I really enjoyed the book and learned a lot from it. I like Clark's self-indulgent but light-touch style which he pulls off about 98% of the time.³ I like his 'extra large cat, Lolo,' of which more later. And I like a lot of the material which, to me as a sociologist, was new. For example, there is a stunning card trick played out on pages 65 and 66 without needing cards which, since sense self-indulgence is a theme here, I will say that I 'got' after an initial 'expletive deleted' and then about three seconds thought. This fascinating section of the book is about how we

² Evan Selinger helped me greatly with this exchange. He pointed to things I should read, pointed out things I had missed in what I had read, and, occasionally, put right my interpretations of the things I had not missed.

³ I think there is too much use of the unpleasant term 'skinbag.'

pay attention only to that which is in focus and that, as indicated, is also going to be the central theme of the criticism that I will develop later.

Brain flexibility and the body

Central to Clark's argument is the claim that the human brain is flexible enough to cope, in a natural way, with all kinds of extensions to its reach and powers. Like Merleau-Ponty's description of a blind man's relationship with his stick, we are equipped to develop prosthesis-like bondings with new things and associated media and this happens quickly and flexibly. The flexibility argument is reinforced by descriptions of psychological experiments. The experiment that I find 'paradigmatic' was carried out by a UCSD professor. He got person 'A' to sit immediately behind person 'B'. A third party, C, then tapped B's nose with A's right hand while simultaneously tapping A's nose with his, C's, left hand. A felt himself tapping a nose with his outstretched right hand while, with the same rhythm, he felt his own nose being tapped. Apparently, after a minute or two, A felt that the tip of his own nose was coextensive with the position of the position of B's nose – he had a two-foot long nose! I can't wait to try it.⁴ There follow several more descriptions of such experiments with equally fascinating results including the famous inverting eye-glass lenses. It seems it does not take long to learn to ride a bike wearing glasses that invert the field of view. By the end of all this I find myself convinced by Clark that my brain is as flexible as he says it is and that I should be ready to treat the new world of electronic prostheses as more of an opportunity than a threat. I am going to try to change from being a natural-born curmudgeon to being an easy-going cyborg (with a few reservations of which more below).

Of course it may be that I am easily persuaded of these arguments because they could be seen to give at least a little support to some of mine. I argue, contrary to the position of Heidegger, Merleau-Ponty, Dreyfus, and various others, that there is a sphere of human life – namely fluent language speaking by individuals -- where the body has no importance and it is the brain that counts; this is the idea of 'interactional expertise' (Collins, 2004, Collins and Evans 2007; Collins et al 2007w). Interactional expertise is

⁴ I have now tried it and, regrettably, it seems to need more perseverance than me and my colleagues could muster.

the ability to speak the language of a community fluently but the crucial point is that it can be done without being able to engage in the bodily activities of the community. A person who engages in the bodily activities has 'contributory expertise.'⁵ The claim is that it is possible in principle, and widely found in practice, that a person can have interactional expertise without contributory expertise. The idea has wide application, though it initially arose out of a puzzle about the abilities of someone like myself, a sociologist of scientific knowledge. I have spent decades immersed in a certain scientific community and have come to speak their technical language fairly fluently. I have even passed a Turing-Test like examination of my ability to speak the technical language – that is, to speak fluently and creatively, not just parrot phrases. But I have never done any experiments or written any papers pertaining to the science in question. Thus, though I cannot 'walk the walk,' I can do more than just 'talk the talk;' I can, as it were, 'walk the talk' (or perhaps it is 'talk the walk').⁶ What we in Cardiff call the 'strong interactional hypothesis' is that someone with maximal interactional expertise and no contributory expertise will be indistinguishable from someone with both kinds of expertise in any test based on language alone. We have shown that this hypothesis is supported in the case of the colour-blind who can pass as colour-perceivers and in my own performance of the language of gravitational wave physics (Collins et al 2006; Collins and Evans 2007; Collins et al 2007w; Giles 2006). The idea of interactional expertise also helps us make sense of the description given by Oliver Sacks of the abilities of 'Madeleine,' a blind and otherwise almost completely disabled person. Sachs says she was a fluent speaker and could talk with verve about activities that she had never herself experienced (Sacks 1985).⁷

⁵ If you have contributory expertise you probably have interactional expertise too but you at least have latent interactional expertise (you would have the ability to speak the language of the community if you were any good at speaking).

⁶ So far as I know, and our experiments seem to reaffirm the point, this ability can be acquired only through immersion in the lived spoken life of the community. It is not to be gained by any amount of reading or dense integration into the internet; fluent language-speaking is a tacit-knowledge laden activity.

⁷ For a debate about the validity and implications of this position see Selinger, Dreyfus and Collins, 2008; Collins et al 2007w).

Following this line of argument it is possible to distinguish between the 'social embodiment thesis' and the 'minimal embodiment thesis.' The social embodiment thesis holds that the language of a human community is partly a function of the typical form of the bodies of its members – a kind of inward-looking Sapir-Whorf hypothesis: just as those who are habitually surrounded by snow might be expected to have many words to describe it, those of us with an upright stance and whose knees bend backwards might be expected to talk of chairs as something to sit on when we want to relax. In contrast, those who walk on all fours and are accustomed to being pushed around with whips and four pronged instruments might be expected to talk (if they could talk), of 'chairs' as sharp and unpleasant instruments of subjugation. That is why Wittgenstein was right when he said, 'if a lion could talk we would not understand what it said to us.' On the other hand, you, as an individual, do not need legs to be able to speak fluently of chairs so long as you have been brought up in chair-speaking society. If a (talking) lion had been snatched from its mother's breast and brought up among us we would understand what it said to us because it would speak our language rather than 'lionese' in spite of its unsuitable body. That claim about the individual as opposed to the group is the minimal embodiment thesis. It states that the bodily form of an individual has a minimal influence on the language the individual comes to speak because the language is determined by the embedding society. Or, to put it another way (which is how it was first set out), an individual needs only a minimal body (ears, mouth, larynx, and maybe something else), to become fluent in the language of the surrounding society. (Each way of describing the minimal embodiment thesis may be a corollary of the other.)

Clark's brain-flexibility argument seems to lend a crumb of support to these theses for the following reasons. If the 'mind' is flexible enough to get itself around any form of body then there is nothing special about the old-fashioned body as we currently know it. If we all began tapping each others' noses instead of our own the common language would soon begin to deal with the nose as something two feet in length. Indeed, our language is changing all the time to take into account the new forms of relationship that the electronic frontier affords us and, on Clark's argument, it could end up very distant from our current language. This seems to accord with the social embodiment thesis.

More important, however, Clark's position seems to give weak support to the minimal embodiment thesis just because it shows the brain is so flexible. As Evan Selinger, with whom I have been arguing the point about the body (eg Selinger, Dreyfus and Collins, 2008), has pointed out to me, Clark's position is interesting in that he does believe in the importance of the body as a driver of the mind but also believes that the mind is capable of dealing with an indefinite variety of bodily forms including virtual bodies. You could say that if you stopped using your physical body and merely embedded the remnants of 'yourself' (your ears, larynx, and brain) in the language of a group of embodied others that would be a way of giving yourself a virtual version of their bodies. Perhaps the minimal embodiment thesis is right because embedding a minimal body (ears, brain, larynx), in language is a way of extending it into a much more complex virtual body. Maybe that is one way to describe what is happening in the experiments that support the strong interactional hypothesis.⁸ And if it is, then it follows that I could learn to speak 'two-foot nose language' even if I had no arms or nose so long as everyone else was speaking it.

Another reason I like Clark's argument is that he is clear that the flexibility he talks about is a property of humans alone. For example, Lolo, the extra large cat, does not have it. Nevertheless, Clark 'misses a trick' when he discusses our interaction with machines and intellibots. To be fair, many of the worries that arise from the treatment of elements of the physical world as extensions of our human selves are discussed in the section of the book entitled 'Alienation' which starts at page 177. But there is a deeper point: Clark readily talks of 'dovetailing' -- extending our abilities in virtual networks, prostheses, and so forth, by fully engaging with them but he does not talk about dovetailing with Lolo. A dovetail joint is a symmetrical joint: if you were to dovetail with Lolo, the Lolo would have to dovetail with you but Lolo can't -- Lolo's brain does not have human-like flexibility. To understand our role in networks we have to be

⁸ Theresa Schilhab (2008; Collins et al 2007w) even argues that some part of this ability might come from the operation of mirror-neurons during the watching of others' bodily activities and Rodrigo Ribeiro (2008; Collins et al 2007w) also stresses the importance of watching activities for the generation of interactional expertise. The case of Madeleine, on the other hand, seems to be a pure case of just listening and talking without any watching.

continually aware of the difference between symmetrical and asymmetrical links and talk very carefully about it. When it comes to machines Clark's talk is careless. In his treatment, anthropomorphism is too easily mistaken for symmetry. Yes, I can interact comfortably and 'transparently' with the extension of my brain on which I am writing this piece but it can't interact comfortably with me; whenever something goes wrong in the interaction it is me who has to repair the problem (Collins 1997). 'Careless talk diminishes lives,' as one might say. There has been far too much careless talk around for the last four decades. A language which treats an easy anthropomorphism as a symmetrical relationship encourages a colossal misunderstanding of the nature of the human; that is why it diminishes us.

Looking back at cognitive psychology

Now to what the book is 'really' telling us. The main thesis is that the brain's interactions with the world about us represents a natural extension of our cognitive powers – that is to say, we have never been isolated thinkers but were always supported by a cognitive scaffolding outside our physical selves. As Clark likes to put it, as thinkers we are not confined to our 'skinbag' – to 'skin and skull.' As a sociologist, what I find odd about all this is that it is presented as something new -- 'a discovery.' But sociologists have always known it. For example, the sociology of knowledge is generally introduced in year one of the undergraduate course or even earlier. The sociology of knowledge is about how what we experience as sound knowledge varies as a function of where we are brought up and the different social forces to which we are exposed.⁹ You can't have a sociology of knowledge unless people can be brought up differently and you can't be brought up differently unless you can be brought up in the first place, and you can't be brought up in the first place unless you are embedded in a wider society. Being brought up is a matter of relationships with other people and other people are, as a first approximation, things outside yourself. So things outside yourself determine your knowledge. So it is completely obvious, at least to the sociologist of knowledge, that the

⁹ The sociology of scientific knowledge (SSK) is the field in which I do most of my work. It aims to show that scientific knowledge is not so different from ordinary knowledge as we once thought.

things we do with our brains have huge amounts to do with what is outside them as well as inside them.

How is it that the role in cognition of what is outside the brain counts as a revelation in cognitive psychology? I can only suggest, following Clark, that it has to do with figure and ground. Human individuals are about the same size as us and so are easy to see and think about. What is more we all have a lot of experience of individuals – we are one. Societies, on the other hand, are big, ethereal, invisible things. On the pages of Clark's book one can see the hard think-work that has been needed to move our vision from the central field of the individual brain just to the very near edge of the surrounding field. The very effort and cleverness that is needed to show us that skin and skull are not a crucial boundary reveals the extent to which cognitive psychology has been concentrating so fixedly only on what is at the centre of its field of vision.

Now, lest this sounds like special pleading from a professional sociologist, and/or determined 'socio-babbler,' let us look briefly at one of the practical results of the failure to grasp the point that the individual cannot be thought about in the absence of the society. One may read through, or sit through, any number of discussions of the promise of neural nets. Neural nets are always on the point of becoming brain-like and, presumably intelligent. Unfortunately, it never seems to occur to the writer/speaker of those expositions with which I am familiar that the individual brain is not the locus of intelligence.¹⁰ In so far as we live in a world of neural nets, the individual brain is just a small part of one of them. The bit of the neural net that is found in the individual human brain has no significant boundary at the skull. Andy Clark is right! Consider language once more. Any neural net that comprises human cognition extends at least as far as all the language users that the individual brain comes into contact with. The material links within any one brain (whatever 'material' means on a sub-microscopic scale), are just a bit of the net with links extending seamlessly to all the other brains, merely switching from 'material' mediation, to mediation by electromagnetic radiation, sound, smell and touch as the skull's boundary is crossed. Draw a neural net for language and the skull represents, at best, an inconsequential dotted line. Any strengthening and weakening that

¹⁰ I understand the Clark has published on the topic of neural nets; he may be an exception. I apologise for not knowing his particular work on this topic.

goes on in the links that exist inside your skull cannot be understood without understanding the strengthening and weakening of links between all the other neurons in everyone else's head which are engaging with your neurons via any of your senses. That is why I say that other people are outside ourselves only as a 'first approximation;' actually we are part of them. An 'intelligent' neural net based on what is inside the skull will, at best, reproduce something like an autistic person, a feral child, or the inhabitant of a Skinner box who can learn no more than a pigeon. Isn't that just plain obvious? How can we possibly build anything that resembles human intelligence unless it is part of the wider neural net – if it cannot be socialized? Yet not one of the intelligent machines that have ever been actually built, or even conceived of, has the capacity to be socialized. The very best we have are machines that 'learn' through variable reinforcement of their output but not since Skinner has it been believed that this is how humans learn – Lolo maybe, but not humans.¹¹ So all the talk we have had of intelligent machines is anthropomorphism and it is no surprise at all that no machine has ever come near passing even so simple a thing as the Turing Test in spite of the riches beyond the dreams of emperors that would follow from success.

None of this is to say that the restricted project of cognitive psychology as we find it is not difficult, fascinating and worthwhile. If we could understand a pigeon, an autistic person, or a feral child to the point that we could model their behaviour in a computer it would be a remarkable achievement. Carried away by my own rhetoric I may not have given due weight to the fact that the individual brain is a specially dense region of nodes and connections and, marvellously, is capable of storing the knowledge that is distributed in the wider net, at least for a while. Sit a language speaker alone in a room and they can still speak the language at the end of the day pretty well as they spoke it at the beginning. It is only in the long term that isolation degrades performance. It is as though the human brain is a well insulated kettle – heat it up and it stays hot for quite a time. Studying hot water is a pretty important thing to do but it must not be forgotten that studying hot water is not the same as studying energy. It is the hype that passes off hot

¹¹ Forgive me if there are pockets of cognitive psychologists who have reverted to a Skinner-like view; this critique comes from a sociologist and is very much painted with a broad brush.

water for energy that is the problem. I suspect there may be cognitive psychologists who find that hype as irritating and dangerous as I do.¹² Cognitive psychology is a great subject – this review indicates some of its more fascinating findings -- but in its current form it won't provide what the boosters say it provides. The intelligent machine is always just around the corner but we will never reach the corner while we think the individual is all there is to intelligence.

Echoes of this sociological critique can be found Clark's argument, in particular that the human cognitive mechanism is not bounded by 'skin and skull.' But Clark, I am arguing, has merely grazed the surface. What is needed, to coin the old cliché, is a gestalt switch. Start to see knowledge as located not in the individual but in society. To refer back to the nice card trick, cognitive psychology sees the Queen of Spades but what does 'Queen of Spades' mean without the rest of the deck and all the players? The central topic of cognitive psychology – or least that part of cognitive psychology interested in modeling the mind by building working computer programs – should become the study of how the single skull-bound brain draws from the collective body of knowledge. The goal of artificial intelligence should become to build a machine that can suck knowledge from the collectivity. Admittedly such a machine would likely share many of the features of the human brain – and that is why cognitive psychology is doing great work -- but to imagine that building something with just those features and none of the sucking power is to miss the point. We must invent the 'culture sucker.'¹³ How far have we got? COG! And yet every other week I read that intelligent robots are just about to take over the world!

¹² Attending a public lecture by Stephen Hawking at the General Relativity 17 conference, I found myself infuriated by the obscurity of his performance and the showmanship of he and his entourage, in which the mass media were complicit. I found the lecture to be a religious event rather than a piece of physics. Cornering some of my colleagues from gravitational wave physics I remonstrated with them 'you see, your subject has its priests and its acolytes too.' I was surprised to find that they were as angry as I was with what they had just witnessed.

¹³ Oddly enough, the rhetoric of the expert systems boom of the mid-1980s pointed in the right direction though the conception of the problem was hopeless (Collins 1990; Collins and Kusch 1998). As soon as one is tempted to make a claim about how a culture sucker might work one should ask: Could a child learn to be a grown-up in such a way?

Conclusion

That's it except for drawing out a few more points by way of summing up. What Clark has done in an attractive and persuasive way is to show that our cognitive powers are not to be understood by looking only inside the skull. He exemplifies his argument by revealing the power of the set of asymmetrical relationships we have with the physical world, though he himself does not seem to have noticed that they are asymmetrical. That he has not noticed the asymmetry is illustrated by his impoverished model of language. For Clark, language is another set of scaffolding, serving the same function as the pencil and paper of the mathematician. Words enable us to make thoughts into physical things that we can then work with in turn so that we can have thoughts about thoughts. It is true that words can work like this but they do much more. Clark does see this in an ill-formed, peripheral vision, kind of way as the occasional, foot-shuffling, reference to 'culture' indicates. To see the bigger picture look at how Clark uses language rather than at what he says about it. To take just one example, look at the work done by Lolo, the extra large cat.¹⁴ Does Clark's mention of Lolo mean only that he has a large cat in his house. No, it means that Clark is (or likes to present himself as), a happy family man. This reading could not be carried off by someone who did not share a culture with Clark. I have to understand deeply the role of the domestic cat in British households and the conventions of academic writing to recognize just how much work is being done by Lolo without anything explicit being said. And Clark, sharing the culture as he does, will understand why I can say with such confidence what extra-large Lolo is doing for his presentation of self and why no current of foreseeable 'bot' will 'get it.'¹⁵

¹⁴ There is a lot more, not least about Clark's intimate relations with his partner, but Lolo will do for our purposes.

¹⁵ Barring retrospective 'feeding by hand' of each such specific piece of information after the fashion of a fixed 'script' – which is not how we learn. (Sad that I felt compelled to write this footnote but experience suggests it is necessary.)

References

- Clark, Andy. (2003). *Natural-Born Cyborgs: Minds, technologies and the future of human intelligence*. New York: Oxford University Press.
- Collins, H. M. (1990). *Artificial Experts: Social Knowledge and Intelligent Machines*. Cambridge, Mass: MIT press.
- Collins, H. M. (1997). "Rat Tale: Sociology's Contribution to the Problem of Human and Machine Cognition in Context." Pp. 293-311 . in *Human and Machine Expertise in Context*, edited by P. J. F. K. M. Feltovich and R. R. Hoffman. Boston Mass: AAAI/MIT press.
- . (2004). "Interactional Expertise as a Third Kind of Knowledge." *Phenomenology and the Cognitive Sciences* 3,125-143.
- . (2008 -- forthcoming). (ed) *Case Studies in Expertise and Experience: Special Issue of Studies in History and Philosophy of Science*, 39, 1, [March].
- Collins, H. M., Robert Evans, Rodrigo Ribeiro, and Martin Hall. (2006). "Experiments with Interactional Expertise." *Studies in History and Philosophy of Science* 37,656-674.
- Collins, Harry and Robert Evans. (2007). *Rethinking Expertise*. Chicago: University of Chicago Press.
- Collins, Harry, Robert Evans and KES. 2007w. ALL @ SEE; The Expertise Network www.cf.ac.uk/socsi/expertise
- Collins, H. M. and M. Kusch. (1998). *The Shape of Actions: What Humans and Machines Can Do*. Cambridge, Mass: MIT Press.
- Giles, Jim. (2006). "Sociologist fools physics judges." *Nature* 442,8.
- Ribeiro, Rodrigo. (2008). "The Role of Interactional Expertise in Interpreting: The Case of Technology Transfer in the Steel Industry." in *Collins (ed) 2008*.
- Sacks, O 1985. *(The Man) Who Mistook his Wife for a Hat*. London: Duckworth.
- Schilhab, Theresa. 2008. "Interactional expertise through the looking glass: A peek at mirror neurons." in *Collins (ed) 2008*.
- Selinger, Evan, Hubert Dreyfus, and Harry Collins. (2008). "Interactional Expertise and Embodiment." in *Collins (ed) 2008*.

The Blind Carpenter: A Reply to Harry Collins by Andy Clark

In "The Cruel Cognitive Psychology" Harry Collins offers an engaging, (mostly) accurate, and profoundly challenging perspective on some central themes from my *Natural-born Cyborgs*. Among our points of agreement we count the rejection of what might loosely be termed body-chauvinistic accounts, according to which the precise details of individual embodiment place strong and immutable constraints on the kinds of knowledge and thought available the embodied agent. Here, Collins endorses what he dubs a Minimal Embodiment Thesis, which stresses the productive 'give' between the details of individual embodiment and the forms of thought and reason made available by the language practices that surround them: practices that may be centered on forms of embodiment not shared by every individual speaker. My own focus hereabouts has been on what I dub the constant *negotiability* of our embodiment itself, stressing the remarkable ease with which brains like ours learn to make the most of new opportunities provided by changing swathes of scaffolding and equipment, be it sensorily, brute-physically, or even *cognitively* enabling. Crucial to the argument of *Natural-born Cyborgs*, then, was the claim that new layers of non-biological scaffolding (pens, papers, software packages and the like) might literally become incorporated into the very mechanisms of (some kinds of) human thought. The mechanisms of your own mind, I wanted to say, are not bound by skin and skull, nor fixed (though they will be constrained) by the wiring diagram of your individual brain. Minds like ours, if I am right, are chameleon hybrids: the emergent products of repeated rough and tumble mergers and coalitions resulting from the opportunistic exploitation of surrounding structures and opportunities.

In describing these iterated mergers and coalitions, I spoke of processes of 'dovetailing' between neural processing and extra-neural resources. And it was here, Collins suggests, that I slipped gently but dramatically off the rails. For dovetails are in one sense symmetrical. Each side, it appears, has altered to interleave the other. Yet in our relations to those enabling swathes of non-biological scaffolding and support, the process of active interleaving looks distinctly lopsided. My pen and paper do not bend the better to

accommodate 'me'. Even my cat, Lolo, though in himself a remarkable piece of kit (sorry), lacks the full flexibility, in this regard, that characterizes the human participant. Asymmetry looms, and we ignore it, Collins fears, at our peril:

"A dovetail joint is a symmetrical joint: if you were to dovetail with Lolo, the Lolo would have to dovetail with you but Lolo can't – Lolo's brain does not have human-like flexibility. To understand our role in networks we have to be continually aware of the difference between symmetrical and asymmetrical links and talk very carefully about it. When it comes to machines Clark's talk is careless. In his treatment, anthropomorphism is too easily mistaken for symmetry. Yes, I can interact comfortably and 'transparently' with the extension of my brain on which I am writing this piece but it can't interact comfortably with me; whenever something goes wrong in the interaction it is me who has to repair the problem (Collins 1997). 'Careless talk diminishes lives,' as one might say. There has been far too much careless talk around for the last four decades. A language which treats an easy anthropomorphism as a symmetrical relationship encourages a colossal misunderstanding of the nature of the human; that is why it diminishes us." Collins (this issue).

I have quoted this at length because it comes so clearly from the heart, makes a point that is in one sense undoubtedly correct and immensely important, yet manages (nonetheless) subtly to miss the mark, at least as far as the argument of the book is concerned. Moreover, this issue lies at the root of much (misplaced) disquiet about the arguments concerning extended cognitive mechanisms (see, for example, some of the worries raised in Rupert (forthcoming)).

To see where such worries go wrong, let's return to the image of the dovetailed joint. Notice immediately that both sides of the joint are deliberately carved to fit each other by a single agent, the carpenter. To whom or what in our story does the carpenter correspond? To no-one at all! In suggesting the idea of a deliberate act of joinery, the image misleads. What we need (and what I had hoped the rest of text would conveniently

imply) is the *blind carpenter*, erstwhile compatriot of the blind watchmaker, and representative of the multiple and multi-time-scale processes by which brains automatically adapt to socio-technological scaffoldings that themselves automatically adapt (by non-biological mechanisms of differential reproduction) to brains. Elements of deliberate design are clearly part of this process. But they do not exhaust it, and the designing itself is seldom done by the very same individual whose brain gets dovetailed to the product in question.

Looked at in this light the asymmetries stressed by Collins, though real enough, can be seen not to threaten the actual argument. It is true that the key micro-locus of plasticity in all this is the individual human brain. It is the brain's great plasticity and thirst for cheap, outsourced labor that drives the distributed engines of socio-technological adaptation and change. It is true too, that subtract those meaty islands of wet organismic plasticity and the whole process grinds to a standstill. No new pens, paper and software packages when the human organisms all dry up and die. But it by no means follows, from the fact that those wet organismic islands are in that way *lopsidedly essential* to all this, that the rest of the hybrid, distributed circuitry is not part of the mechanistic base for specific episodes of cognitive processing. Similarly, as Dave Chalmers (personal communication) has recently pointed out, the activity of my visual cortex is not essential for my being the agent that I am. Were my visual cortex damaged, I would persist. Nonetheless, while up and running, it functions as part of the physical machinery upon which 'I' supervene. My relation to my own visual cortex is thus itself lopsided, but this in no way renders it false to speak of the visual cortex as helping to realize some of my mental activity.

We should bear this kind of case in mind when Collins (see previously quoted passage) notes that when his interactions with props and tools go wrong it is 'he' who has to repair them. For the same may be said when aspects of our own bio-memory start to become unreliable and 'we' shift towards alternative means of storage and retrieval. Once again, the apparent lopsidedness (I have to take steps to offset the loss of bio-memory

functioning) does not threaten the claim that, prior to the loss, those bio-memory resources were realizing *my* cognitive activities.

In general, the appearance of lopsideness arises, I suspect, largely because we are unused to thinking of our brains as themselves not a single indivisible unity ('me') but simply another collection of mechanisms upon which the agent currently supervenes. The agent is, to use a term from Hurley (1998), a kind of singularity that may be realized, at different times, by different coalitions of resources.

If this seems a little murky that is because it is, in fact, *extremely* murky. But it is a good murk. It is the kind of murk we have to navigate, and soon, if we are to appreciate ourselves both as human agents and as shifting flows and eddies in the flux of mechanical cause and effect. What Collins does, with elegance and penetrating wit, is to push just hard enough at the story I tell to force it to reveal its true, though barely (if at all) realized, intent. For I think that we need to do more than simply embrace the profound role of socio-technological scaffolding in the construction of mind, thought and reason. To go that far may be simply, as one writer recently put it, to apply a "lexical band-aid" to "a 350 year old wound generated and kept suppurating by a schizoid metaphysics" (Maxine Sheets-Johnstone (1999 p.275)). To really make progress we need to address the suppurating metaphysics itself. We need to go beyond the uneasy superposition of intuitive dualism and organismic-individualism that underwrites much present debate: to know ourselves in some way that does justice both to our experience of persisting agency and to its shifting and heterogeneous material base.

References

Collins, Harry (this issue) "The Cruel Cognitive Psychology"

Rupert, R (forthcoming-b) Innateness and the Situated Mind, in P. Robbins and M. Aydede (eds.), *The Cambridge Handbook of Situated Cognition* (Cambridge UP)

Sheets-Johnstone, M. (1999). "Emotion and Movement: A Beginning Empirical-Phenomenological Analysis of Their Relationship." *Journal of Consciousness Studies* 6(11-12): 259-277.

The Cruel Harry Collins, by Jeff Shrager

In his essay, “The Cruel Cognitive Psychology”, Harry Collins’ takes Andy Clark’s argument of brain flexibility in ‘Natural Born Cyborgs’ as evidence “there is a sphere of human life – namely fluent language speaking by individuals -- where the body has no importance and it is the brain that counts; this is the idea of ‘interactional expertise’”. Whether or not it pans out scientifically, this is a deep and important hypotheses, and I agree that Clark’s gathering of evidence lends support to it. That, declares Collins, is an interesting story that he finds barely submerged in Natural Born Cyborgs. But Collins dives even deeper to find what Clark is “really” telling us – the submarine story, so to speak – that “as thinkers we are not confined to our ‘skinbag’ – to ‘skin and skull.’” That that “things we do with our brains have huge amounts to do with what is outside them as well as inside them.” For sociologists of knowledge like Collins, this is of course mother’s milk. But Collins then spends the rest of his essay berating cognitive psychology for missing it: “Any neural net that comprises human cognition extends at least as far as all the language users that the individual brain comes into contact with.”¹⁶ “How is it that the role in cognition of what is outside the brain counts as a revelation in cognitive psychology?”

Collins is likely referring here to one of two general theories: Either he is (a) imagining that all our brains have to be somehow physically and in real-time connected up in order to anyone to think anything, or (b) imagining a temporally-extended analysis of cognition that reaches back to one’s infancy. Let’s leave aside the first of these as perhaps an appropriate discussion for another time, with more allotted space-time. As for the second theory, a historiologically-intensive analysis of the history of language and thought could be useful, although it would probably be hard to pull-off. (Where does it end? Do we need to go back to the obelisk that came down from Jupiter and gave the early hominids

¹⁶ Neural nets stand in for Collins throughout as the paradigmatic cognitive model; although I don’t think that they are in fact paradigmatic, I think that Collins would assert that his argument holds for whatever sort of model a cognitive psychologist would wish to offer, so I’m happy for the moment to let neural nets play this role.

their first words?) But, just as the importance of one's culture on thought is old news to sociologists, it is old news as well to psychologists, and even older developmental psychologists than to Harry. For example, in the early part of the 20th century the Russian developmental psychologist Lev Vygotsky developed a fundamentally socio-cultural theory of cognition wherein what he called "cultural mediation" and the "internalization of social speech" were considered fundamental aspects of the origins of cognition; and this set of theories is quite well respected by psychologists to this day.¹⁷

Neither is it news to computational developmental psychologists¹⁸ that one's cultural surround is critical to cognition. In fact, Clark's book is based largely upon this premise:

"Our distinctive mathematical prowess depends on a complex web of biological, cultural, and technological contributions. First, the biological brain commands an approximate sense of simple numbers. Second, specific cultures have coined and passed on specific number words and labels, including key innovations such as words for zero and infinity. Third, the cultural practice of enforcing simple rote-learning regimes (mathematical tables and so forth) added another element to the matrix. Finally, mix in the novel resource of pen and paper, and PRESTO! Our culturally enhanced biological brains can begin to tackle and solve ever-more-complex problems, eventually scaling mathematical heights that our unaided biological brains (of our stripe) could never have hoped to conquer."

And Clark didn't just make this up, he cites numerous papers by cognitive psychologists, developmental psychologists, and computational developmental psychologists. Indeed,

¹⁷ Vygotsky's ideas, and more generally the importance of inculcation in cognition, were reintroduced into the mainstream of developmental and cognitive psychology in the 1980s by Wertsch, and others (Wertsch, J. V. (1985). *Vygotsky and the Social Formation of Mind*, Harvard University Press, Cambridge, Mass., and London), and their importance is now widely discussed.

¹⁸ I would probably be considered a "contributory" computational developmental psychologist, based upon Collins' and Evans's periodic table of expertise where a "contributory expert" can undertake an activity with competence. Or, at least, I've managed to convince a few peer reviewers of my limited competence.

Even deep connectionist brain modelers, working even before the ones cited by Clark, and whose work those later researchers based their work upon, understood the central role of culture in brain development:

"The mammalian cortex is organized during development through a combination of endogenous and exogenous influences including genetic restructuring, subcortical influence, maturational timing, and the information structure of the organism's early environment." [...] "The importance of the present theory with respect to the development of human cognition has not escaped our attention. One of the principle aspects of human development that distinguishes us from our nearest primate neighbors may be changes in the timing of development. Among primates, the human brain remains relatively plastic until late in the post-natal period, whereas the brains of our closest relatives are more completely formed by birth. Thus, humans have greater access to the complex experiences afforded by being out in the world during the most formative period of brain development. This may enable us to "tune in" to rapidly changing features of the environment that other animals, more constrained by evolutionary structuring of their brains, are unable to take account of. To the extent that these changes are reflected in similar timing changes in the hypothesized wave of cortical development, they may also lead to changes in the extraction of information from the stimulus environment during the early self-organization of the brain. Of course, along with this flexibility comes an extended period of immaturity, during which we are dependent upon our caretakers and our community for both support and training. *Our theory therefore sees the coevolution of culture and cognition as a fundamental condition of human evolution.*"¹⁹ (emphasis added)

And as early as a decade before that, non-connectionist computational developmental psychologists were training their models with inputs whose statistical distribution

¹⁹ Shrager, J. & Johnson, M. H. (1996). Factors influencing the emergence of function in a simple cortical network. *Neural Networks*, 9(6), 1119-1129.

mirrored that observed in parent-child interactions. This practice has only grown in importance since that time, and the analysis of what the social environment provides the child, and of how the child uses that environment, have become increasingly sophisticated.

So if Collins' volley regarding the (non)revelation of the importance of the social in cognitive psychology is a shot across the wrong bow, at whom should he have been aiming? Collins gives us some hints, complaining that 'it hasn't worked' yet. ("How far have we got? COG!"). Here I take it that the "it" he's talking about is the program of building intelligent machines. Although cognitive psychologists aren't trying to build intelligent machines, I'm happy to stand up in defense of AI engineers as well.²⁰ Not having succeeded at something is not a very strong argument that its theoretical basis is wrong, just as our failure as yet to create a quantum computer is no indictment of quantum mechanics. And I haven't seen sociologists building artificial societies that work any better than AI folks' poorly-working intelligent machines! Neither he nor we have HAL! So when Collins says "What is needed, to coin the old cliché, is a gestalt switch. Start to see knowledge as located not in the individual but in society." it rings hollow; sociologists of knowledge say this a lot but I've never actually seen it cashed out in a detailed analysis of a system (working or not), except to complain that this or that computer program ("COG!") isn't "really" intelligent. Okay, I agree that COG isn't HAL, but under what analysis isn't this "real intelligence"? Maybe it's only the equivalent of a baby of few months old. Is a baby "really" intelligent? A baby's brain is a whole lot more complex than anything we can build in any current computer, so it isn't too surprising that we don't even have "real" babies yet. Again, applied failure is a very weak argument against theory.²¹

²⁰ I think that I would count as a "contributor" expert in some areas of AI as well, so I suppose it's my duty to stand in their defense!

²¹ And Collins' theory that "Human individuals are about the same size as us and so are easy to see and think about." seems an odd thing to say, esp. to a developmental psychologist; babies aren't "about the same size as us." Do developmental (cognitive) psychologists somehow need less of a switch? Are developmental (cognitive)

Speaking of HAL, the super secretive, super intelligent computer in another Clarke's (Arthur C. – spelled with an extra 'e') *2001*, there is in Collins' essay some evidence that who he's really complaining about is science fiction writers, not cognitive psychologists. He complains: "And yet every other week I read that intelligent robots are just about to take over the world!?" Where? In science fiction thrillers about super secretive, super intelligent computers and apes getting language from ancient Jovian obelisks? And even "Arthur C-with an e" already knew what Collins claims to be somehow the mystical rite of sociologists; In the film adaptation of *2001*, which appeared in 1968 – barely ten years after the birth of the modern computer! – there is a brilliant scene where Dave Bowman turns HAL off by pulling out what are presumably parts of HAL's neural net. As HAL loses his memory and begins to deteriorate, his speech slows and slurs and he reverts to his child-like self; who can forget HAL's words, and especially the song:

HAL: "I'm afraid" "I'm afraid, Dave" "Dave, my mind is going." "I can feel it."
"I can feel it." "My mind is going." "There is no question about it." "I can feel it."
"I can feel it." "I can feel it." (slowing speech) "I'm a...fraid." (very slowly, as if a child, HAL) "Good afternoon, gentlemen. I am a HAL 9000 computer. I became operational at the H.A.L. plant in Urbana, Illinois, on the 12th of January 1992. My instructor was Mr. Langley, and he taught me to sing a song. If you'd like to hear it, I can sing it for you."

Dave: "Yes, I'd like to hear it, Hal. Sing it for me."

Hal. "It's called...Daisy. (HAL's speech slows progressively, eventually falling into silence) Daisy, Daisy, give me your answer do. I'm half crazy, all for the love of you. It won't be a stylish marriage, I can't afford a carriage, but you'll look sweet upon the seat of a bicycle built for two..."

psychologists not really cognitive psychologists because we (developmentalists) study how babies learn to be humans through (among other things) cultural embedding?

So, apparently Harry and I, and both Andy and Arthur C. Clark(e) – sociologists, psychologists, and even science fiction writers – agree that interaction with one's culture is a critical substrate for cognition. Studying cognition, or building AIs, (or writing science fiction stories) does not, so far as I can tell, in any way exclude cultural embedding. Far from needing gestalt switches, all that we psychologists claim is that the thing inside our skin-and-skull, no matter how you analyze it, is a pretty interesting machine. We just want to know in how it works. And Collins does too!

Symmetry and Socialisation: Reply to Clark and Shrager by Harry Collins

I enjoyed Andy Clark's elegant response to my criticism of his Natural-Born Cyborgs but I could not understand all of it. For example, I do not see why the invocation of a blind carpenter making the dovetail joints makes any difference. I was never in any doubt about the work done by the blind carpenter but I can't see why this renders the relative roles of human agents and material scaffolding any less asymmetrical. I do see the interest in the remark that the asymmetrical relationship involves a somewhat murky agent who has an equally asymmetrical relationship with the physical constituents of his or her own brain. I also see still more clearly that, therefore, this murky agent, can be physically constituted out of a various different collections of bits and pieces and, still be what it is – something like soup. I like the point especially since it allows me to replace much of the 'long pig' with linguistic fluency and still have a nourishing agency broth.

It also seems true to me that we may need a new metaphysics to deal with the agent but I do not think the problem with the metaphysics is what most people think it is. It seems to me that for the last few decades I keep hearing complaints about dualism and the problem with our metaphysics being its sharp divide between the material and the human, or some such. Indeed the words 'false dualism' seems to have achieved the status of a kind of 'mantra,' whose very utterance is capable of bringing an end to debate. As far as I know, however, one of the conditions that must be satisfied by a metaphysics is not to be at variance with the physics supported by it. In this case the physics is the asymmetry between the murky human agent and the material: the fact of this dualism smacks us in the face every day. It is not, then, false dualism that is the problem but unexplained dualism. We need a metaphysics of murky agents but it will not succeed by wishing away the distinction between human agency and the material. This is not the place to go into the arguments, but I am sure the new metaphysics must start by exploring the peculiar nature of the social relations that tie the murky agents together and give them their qualities.

On the matter of cognitive psychology, it looks like I owe apologies all round. This whole interchange has been a learning experience for me. In truth, I have never been quite sure what cognitive psychology is but if it is what Jeff Shrager says it is then my future (if I have one) lies in cooperation with cognitive psychologists not conflict with them. An apology to Andy Clark is needed because I did not point out, as Shrager does, that his remarks about the sources of his own argument show that he too is well aware of other sources for the idea of the extended mind.

It looks like what I have been doing is some poor sociology of academic disciplines. I am still puzzled by the 'revelatory' tone in respect of Clark's claims about the way the mind extends itself beyond skin and skull in respect of material artifacts since his very own work, I now realize, reveals that for a very long time various kinds of psychologist have known that the brain reaches out all over the place. So I am going to pass just a bit of the blame to Clark. Nevertheless, what I have been doing is reading far too much into some rhetorical flourishes – I was just taking it that Clark's way of presenting his ideas implied that the idea that the brain was not 'on its own' was news to lots of people (sociologists apart) whereas I should have seen from the book itself that it was not news at all except in respect of the material dimension of what is outside the skull.

Now, in the last paragraph I say 'sociologists apart' but those two little words again exemplify bad sociology. In fact, over the last two decades, the majority of sociologists who work in 'my neck of the woods,' if it is still my neck of the woods, have worked with views far more sympathetic to Clark's way of seeing the world than to my way. Astonishingly, sociologists of science, such as that very large numbers who align themselves with Bruno Latour's ideas, allow for no ontological distinction between the human and the material.²² My charge to Clark that he misses the asymmetry of the relationship between the agent and the entities from which the agent is composed, would apply equally to this group and they would find his complaint about suppurating metaphysics exactly to their taste – indeed, they could have made it themselves. In future

²² A critique of Latour's work along these lines can be found in Collins and Yearley, 1992.

I will try to heed the warning that talking of this or that kind of 'ists' (psychologists, sociologists, computer scientists), is likely to conflate a lot of diverse opinions.

I can't resist responding to Jeff Shrager's volley of quotes with one originating from his own university. The Stanford News for April 20 2006 reports that a so-called 'singularity summit will be held on the campus. It goes on to explain:

"The Singularity will be a future period during which the pace of technological change will be so rapid, its impact so deep, that human life will be irreversibly transformed," said Ray Kurzweil, keynote speaker and author of the best-selling *The Singularity Is Near: When Humans Transcend Biology* (Viking, 2005). "Based on models of technology development that I've used to forecast technological change successfully for more than 25 years, I believe computers will pass the Turing Test by 2029, and by the 2040s our civilization will be billions of times more intelligent."

I know Jeff is going to say that Kurzweil is more of a science fiction writer than anything else but why cannot we cognitive psychologists and sociologists shut these people up. There are many more of them, as a bit of Googling will show, not least the same Kevin Warwick who is discussed by Clark in his book. Why can't we make the idea that the human is essentially a social animal so salient that it becomes impossible to put forward these outrageous claims without adding 'so long as we solve the problem of socialization.' And COG remains relevant for the same reason. Here is Rodney Brooks, one of the most successful figures in the AI world, working at what is probably the world's leading AI laboratory, building something one step beyond a talking doll and calling it socialization. It brings to mind the Pacific Islanders and their quest for 'Cargo.' In May 1994 I gave a paper to a major British cognitive science department with a strong reputation in AI research in which I argued for the centrality of socialization to the problem of machine intelligence. I was told that I should go and look at COG. When I said that I knew enough about COG to know that looking at it was a waste of time it was intimated that I was some kind of Luddite. Surely this really does indicate that someone

is not doing their job, be it the AI community itself or the people with which it speaks, including ourselves. Or maybe I am doing bad sociology again, I've certainly been shouting my head off for a long time without making any impact and maybe everyone else is shouting too. What we do know is how hard it is to change the minds of a group of academics which is doing well in terms of money and/or public recognition.

Sociologists and psychologists are 'natural born enemies' because they are born fighting over the same terrain. Most of my fights with psychologists have been exactly about the topic under discussion here – the extent to which you can make sense of individuals without seeing them as essentially part of some murky metaphysical thing called the social collectivity. Not to lose the frisson entirely let me say that there is more to the collectivity than Shrager's 'initial upbringing.' There is also the deeper point that no new concept makes sense until it comes to belong to the collectivity. That is the difference between the scientific genius and the crank, the daub and the avantgarde masterpiece, the leader and the blunderer. Our conceptual life cannot be understood in the absence of the collectivity. This is not to say that if you lock me up in a room isolated from my community in the morning I won't any longer be an English speaker by the end of the afternoon. But I won't have contributed to the adding of anything new to the English language in that time and if you lock me up for a lot longer, when I get back to society I will find that English is no longer what it was when I left it.²³

That said, there are few finer experiences in academic life than to find that some group from which you have felt separated because of tribal allegiances are really thinking along the same lines. I just hope that we can all be a bit more vocal in the future about pointing out that socialization is one of the central problems of AI -- I think it is the central problem -- and that no computer will act intelligent until we have learned how to bring it up like a baby, and make it to interact with those around it like an adult. And interacting with those around it like an adult means putting as much into the relationship as it takes out. I don't care if cognitive scientists and computer scientists don't know

²³ For more on the 'episodic' nature of social embedding see, for example, Kusch 2002, p177

how to begin this project, only that they see that it is a problem that must be solved if real progress is to be made in either understanding or reproducing human thought and action.

Individual and collective, by Jeff Shrager

Collins claims (above), "...no computer will act intelligent until we have learned how to bring it up like a baby, and make it interact with those around it like an adult. And interacting with those around it like an adult means putting as much into the relationship as it takes out."

Given this sentiment, I think that we (or, at least Harry and I) are now in complete agreement and there is no need for a further response from me! **(EMS)**

Agents and asymmetry, by Andy Clark

Collins certainly makes you think. His reply is typically elegant, forceful and truthful (a word I'd seldom use in academic context, but that seems apt in this instance).

I guess the point for me is that I don't want to deny the asymmetries Collins (very properly) highlights. The point is murkier than that. I want to allow the agents *themselves* to be constituted by collections of elements whose internal relations may indeed be asymmetric. In that way, the physics and the metaphysics would still fit together, as Collins rightly insists they must. To take a mundane case, my index finger and I have just such an asymmetric relation: I can persist without it but it cannot persist without me. I can help repair it (with a plaster) but it (working alone, as it were) can never repair me...etc, etc.

My claim is that once we look to the parts rather than the whole, it is "fingers all the way down": that where cognition and mind are concerned, we murky agents do not even cleanly stop where we find the bounds of skin and skull. I suspect this is something Collins might even agree with, while still laying more stress on the asymmetries than I do. In a sense, our relationship with language (not with the agents who produce it, but with language considered as a mass of vibrations, inscriptions etc) is also asymmetric...but that doesn't seem to preclude its laying a key cognitive role in his own model. Does this help clarify the issues at all?

Squarks and language, by Harry Collins

I'm still learning. I now see more clearly than ever before that the brain, as we know it, which I have always vaguely thought of as something special, is not much different from the lead in the mathematician's pencil: brain substance, as we know it, may have special properties but only in the way that the lead has the property of making marks on paper. Nevertheless, I still do not understand the determination on the part of Clark and so many others to eliminate the unexplained part of the self (which is to say that my understanding is in terms of intellectual fashion rather than ineluctable logic.) Assuming 'constituted' has an ontological force to it I don't see why Clark "want[s] to allow the agents *themselves* to be constituted by collections of elements."

As far as I can see the asymmetry that we agree exists between the agent and its currently palpable elements leaves the agent with a special role. I confess that I have difficulty with the idea that the special role is simply something that 'emerges' from the way the palpable elements are assembled but maybe that's my problem and I should simply think harder. But how hard should I think? Maybe one can think too hard! Even if the spirit of scientific investigation abjures the invocation of a soul or its equivalent we still know very little of the material constituents of the universe so why should I be determined to explain everything in terms of our current notion of the material elements? To base one's understanding of the human agent on the notion that there is nothing more to know about its material constituents than pencils, muscles, and molecules may well place too great a burden on the couple of seconds science that complete the metaphorical year of human history. Even as I write, physicists are arguing about the relationship between consciousness and quantum uncertainty (those darned cats taking center stage as usual). [AC writes: Regarding pencils, muscles and molecules, it is the functional roles they play that matter, so it is not as staunchly physical-reductionist a project as this makes it sound. HMC responds: maybe that leaves a role for squarks – see below]

Let me hypothesise, then, that the self is a material thing but it is not made out of the elements of brain substance as we know it, even with the addition of arms, legs, bodies, pencil, papers and screwdrivers, but is part of an '11-dimensional, quantum-entangled, tachyonic, dark energy driven, ensemble of chaotic wormholes,' which I'll call

a `squark'.²⁴ Incidentally, one of the things about a squark is that it links all human minds together via quantum entanglement and wormholes so that when we think we are speaking for ourselves we are merely participating in the wholeness of the extended mind of humankind, being particularly strongly connected to those who speak the same natural language. It would be presumptuous to call this a new material property of the brain because that would be to assign a physical location to it that may or may not be appropriate (and hence the references to `the brain as we know it'). A squark, in case it isn't obvious, is not something that I really want to posit – I am not, as many physicists do, demanding a paradigm shift in physics in the face of the intractable problems of the measurement problem in quantum theory, mind over matter (moving my little finger, the placebo effect), and so forth. I am not a physicist. A squark is simply a place-holder for all the stuff we don't know about, the current impenetrability of which becomes evident even to non-physicists if you hang around physicists long enough or just read the New Scientist.

What kind of thing is a squark (and the things it holds a place for)? Dennett (1992) makes the useful distinction between abstract objects, such as the concept of `centre of gravity' -- something useful for doing calculations but merely as an analytic consequence of our way of thinking about mass – and inferred entities, such as (at one time), the atom. You cannot see a centre of gravity because it is only a calculative convenience; you cannot (could not) see an atom but that was no reason not to treat it as real. You also cannot see a squark, but that is not because it is an abstract idea like a centre of gravity, it is because our science and technology is so primitive. All we know about it is what we can infer from a few of its effects on our lives (the day-to-physics of life that must be reconciled with the metaphysics). A squark is not like a centre of gravity, then, but like gravity itself – something that Newton thought up to explain what he saw happening around him even though he could not touch it, see it, or smell it and,

²⁴ Dennett (1992) writes: `Brain research may permit us to make some more fine-grained localizations, but the capacity to achieve *some* fine-grained localization does not give one grounds for supposing that the process of localization can continue indefinitely and that the day will finally come when we can say, "That cell there, right in the middle of hippocampus (or wherever)--that's the self!"' But why that cell? Why not that squark?

worse, it seemed to be 'action at a distance.' Yet gravity is now enshrined as one of the four basic forces and a squark has this potential.

Squarks, of course, may not turn out to be the key to all this stuff happening around us. It may turn out to be quarks, or querps. Or it may turn out that, as Dennett and Clark argue, there is nothing extra needed. Or it might turn out to be an emergent property of the material world we know but it emerges at the level of collectivities of brains, the linking medium not being quantum entanglement, wormholes, or anything like that, but just sound, sight, touch, and smell (as argued earlier, once you start assembling things into an ensemble there is no reason to stop at individual brains). It might be any of these – it is just that we don't know for sure). [AC writes: I am not opposed to squarkiness, just don't yet see any compelling reason to start looking out there for the nature of mind. Are we going this way just because the other story seems so hard to swallow? Squarkiness is pretty indigestible too! HMC responds: One doesn't have to digest squarks per se, just the idea that there is lot out there we still don't understand. But we don't want to swallow too hard on what we do understand.]

Is the attribution of squarkiness to the brain a reinvention of the 'Cartesian Theatre in a thin disguise? The squark idea does indeed imply some as yet unknown properties of agency that 'lie behind' what we can currently see happening in the brain and may well give it direction and its unique ability to 'repair' broken bits of the things it interacts with but it does more than fill a logical lacuna, it might one day be talked of like atoms or gravity. Asquark is certainly not a homunculus -- it is material stuff quite different to a little human. It does not seem to me that one can ban all new discoveries about how agency might work by intoning 'Cartesian Theatre' any more than intoning 'heretic.'

Another thing that a squark is not like is 'the dormative property of opium' -- as in, 'opium makes you sleep because it has a dormative property.' A squark is not a definition in terms but something that might or might not turn out to be real. Finally, the squark hypothesis is not like, say, the soul, or the intelligent design hypothesis. Both of those hypotheses come from the domain of religion rather than science and both are damaging to science because they put an end to scientific inquiry: they say, for example, 'don't search any further for the causes of the development of this organ it was a matter

of intelligent design,' or don't search any further for an explanation of the asymmetrical relationship between the human agent and the pencil, it is a matter of the soul.' The squark hypothesis says 'keep on searching – agency may be a matter of squarks, or maybe not, and one day we might find out.'

Oddly, the position taken by Dennett and, perhaps, Clark, seems to bear more of a relationship with the intelligent design idea, or the idea of the soul, than it does with the squark hypothesis. This is because it implies: 'Don't search any further, we know the answer already.' It implies: 'Everything you need to know about agency, the self, etc. is already there in the material world as we know it.' It demands: 'Just assemble the already known elements and nothing more.' It is, in short, something like a religion of the material world as we know it now. Of course, it is not a religion in the sense of invoking a realm of the spiritual as opposed to the material but it is like a religion in the way it sets boundaries to inquiry in a priori way, bans certain kinds of thought, and discourages heresy with accusations such as 'false dichotomy!' In the face of the-material-world-as-we-know-it religion, I want to insist that there may be squarky stuff out there. [AC writes: I would categorically deny that there is any a priori boundary setting going on...I just await good reason to go looking among the squarks, and also, some sound methodology for doing so (either of these would justify some explorations,) HMC responds: I can't supply the methodology but I think I can supply the good reasons – asymmetry and collectivities. Kuhn's term 'anomaly' may be germane here. Kuhn uses the term to mean two things: anomalies(1) are irritations inconsistencies that one should brush over, or leave to the future, if one wants to get on with things; anomalies(2) are obstinate problems that indicate that there is something wrong with one's current view of the world. I am trying to raise the status of 'asymmetry' and 'role of the collectivity' from the anomaly(1), which is what they are at best in the Dennett/Clark world view, to anomaly(2).]

What the religion has in its favour is the useful rule not to multiply entities beyond what is necessary. Nevertheless, we must always bear in mind the counter-apophism, seemingly a paraphrase of Einstein: 'Theories should be as simple as possible but no simpler.' Were it not that the everyday physics keeps slapping us in the face – were there are everyday phenomena that one can only reconcile with the material world

as we know it by arguing like a zealot – then the squark hypothesis would not have a chance. As it is, however, it seems more than likely that we do need something extra.²⁵ The real danger is that the zeal for parsimony – thinking too hard -- causes us to ignore phenomena that do not fit the doctrine, such as the power of the collective on our lives and the asymmetry of the whole and the parts [anomaly(1) and anomaly(2) HMC.]. The danger, to revert to the psychology of perception motif, is that these things are determinedly or unconsciously relegated to the periphery when they should be in focus.

Clark's final remark about language makes me realize that I have been making mistakes again. I have been far too lax in talking about language as a prosthesis, 'replacement for the long pig,' and so forth. Language, according to my account, is itself something that belongs in the realm of agency – the agency of the collective. Unlike brain substance as we know it and pencils, it's a bit squarky. Language is not just a tool, like a hammer, the asymmetry of whose relationship to the agent allows it to be replaced by a heavy stone. You can't replace language with something else that would do the same job because the language is integral with the life of the collective agent (as manifested in interactional expertise). [AC writes: I don't quite see this...though I do agree that anything that would do that very job would count as some kind of language...is that the idea? HMC responds: Well, we might be able to replace language with something else 'with language-like properties,' but I cannot imagine what it might be like. That is to say, it would be more like a squark than a stone.]

Language, then, is on the wrong side of the asymmetry relationship to be a prosthesis. Language is not part of a virtual body, but it can be a creator of virtual bodies. The form of the virtual bodies it may create is determined by the collectivity's (not the individual's) interactions with the world. As I have learned from this exchange, 'the world' referred to here includes the body and the material objects with which it interacts. As Clark points out, there is enough flexibility for this language to respond to radically changed material circumstances.

²⁵ I am not the only one to believe this – see, for example, various works on the mind by physicist, Roger Penrose. A well-known principle in sociology of scientific knowledge can be summed up by the phrase: 'distance lends enchantment.' Physicists inhabit a far more mysterious physical world than those distant from physics.

What we need, and forgive me for banging on, is to find out how the individual agents interact with the collective agent. We do need to 'keep the collectivity in mind.' It may be that we already know enough to work this out, given enough zeal or it may be that there is more to be discovered.

References

Dennett, Daniel C. (1992) The Self as a Center of Narrative Gravity. In: F. Kessel, P. Cole and D. Johnson (eds.) *Self and Consciousness: Multiple Perspectives*. Hillsdale, NJ: Erlbaum.

Kusch, Martin, (2002), *Knowledge by Agreement*, Oxford: Oxford University Press