

Cognitive Change and Material Culture: a Distributed Perspective

Michael Wheeler

1. The Problem of Ancient Minds: from Neuro-Centrism to Distributed Cognition

Whichever way one looks at it, studying ancient minds is a challenge. This is largely because the main investigative strategies that have proven to be so successful in unravelling the secrets of modern minds are rendered unavailable. Long-dead subjects cannot sign up for reaction time laboratory experiments and decomposed brains aren't suitable for neuroimaging. Call this epistemological quandary 'the problem of ancient minds'. Given that the questions addressed by the present volume concern the existence and character of changes in consciousness and cognition at the Neolithic site of Çatalhöyük, there seems little doubt that the problem of ancient minds will be prowling the pages of the various chapters, just itching to make a nuisance of itself.

Of course, human ingenuity knows no bounds and various canny tactics have been deployed to make progress against the problem, in its general form. For example, in one branch of so-called neuroarchaeology, living human beings have their brains scanned while performing the same tasks as were performed by our ancient ancestors (e.g. making certain tools). On the strength of pre-existing hypotheses from cognitive neuroscience about the psychological functions of the neural areas thereby activated, these experiments generate inferences about which cognitive capacities are required for the execution of the tasks in question, and thereby about the cognitive capacities that must have been present in ancient minds (for a paradigmatic example of this research programme in action, see e.g. Putt et al. 2017).

One way of thinking about what's going on in such neuroarchaeological experiments is to note that the physical machinery whose contribution forms the ultimate target of the investigation is (was) located inside the skull and skin of the cognitive agent. Of course, there are relevant elements located outside the skull and skin too – e.g. tools, artefacts and the other constituents of material culture – and these form part of the traditional, routine business of archaeological research. On the face of things, however, the difficulty, when one's topic is ancient *minds*, is that such external elements are not themselves constituents of our ancestors' psychological machinery. They are 'merely' the products of that machinery and/or items with which the ancient inner psychological machinery of interest once interacted. So, for the archaeologist interested in ancient minds, the struggle, it might seem, is to find reliable ways of using what evidence and resources we do have to uncover the character and structure of the missing ancient inner, *because that's where the minds in question were*.

For some (although not all) readers, the idea that ancient minds were located inside ancient heads will seem to follow pretty much directly from the undeniable thought that ancient

brains were located inside ancient heads, on the grounds that, to the extent that it makes sense to talk about minds being located at all, minds are always located where brains are. Fuelled most recently by the excitement surrounding contemporary brain-imaging techniques, there is no doubt that, right now, this neuro-centric conception of mind is the default view in cognitive psychology, developmental psychology, cognitive-science-friendly philosophy, and indeed in many other intellectual disciplines that concern themselves in one way or another with psychological phenomena. Importantly, in the present context, the reach of such neuro-centrism extends not only to neuroarchaeology, but also to cognitive archaeology in general, understood as the broader endeavour to study the minds of our human ancestors, based upon the surviving material archaeological remains and by drawing on the theories and concepts deployed by contemporary cognitive science. (For just one pioneering example of the broader endeavour, see Mithen 1996.)

Notwithstanding the overall dominance of neuro-centrism, however, the fact remains that, even within cognitive science itself, there is an increasingly prominent alternative. According to the *distributed* view (often traced first to Hutchins 1995), cognition (understood liberally as encompassing mind, thought, intelligence, reasoning, emotions, feelings and experience – in short, the psychological) is, at least sometimes, in some way, spread out over the brain, the non-neural body and, in many paradigm cases, an environment consisting of objects, tools, other artefacts, texts, individuals and/or social/institutional structures. The phrase ‘in some way’ introduces a deliberate vagueness into the specification of the distributed view. As we shall see later, how one fills in the details of the ‘spreading’ at issue will determine whether the distributed view is, at a fundamental (one might even say ‘metaphysical’) level, an illuminating modulation of the neuro-centric view or an altogether more surprising position. For now, however, what’s important is this: in both of the forms in which we shall consider it below, the distributed perspective highlights the important roles that external elements routinely play in experience, thought, reason and so on, but it does so without seeking to marginalize the manifest importance of the brain in the generation of these phenomena. Rather, it aims to place proper emphasis on the point that, to understand what the brain actually does, one needs to take account of the subtle, complex and often surprising ways in which that organ is enmeshed with non-neural bodily and environmental factors.

To give the flavour of distributed cognition, here’s a brief description of what is now a canonical example. Early modern theatre companies performed an astonishing number of plays (as many as six different plays a week), with relatively infrequent repetition, very little group rehearsal, each actor playing multiple roles, and in the face of mounting a new play roughly every fortnight. If we imagine a scenario in which each actor was required to store each of these plays in his brain, it seems that we would be forced to conclude that early modern actors possessed super-human organic memory capacities. This seems unlikely, so how did they do it? Adopting a distributed perspective, Tribble (2005, 2011) argues that a number of tricks and ploys resulted in the seemingly prohibitive information processing required being spread out over the individual actor and the physical and social environments of the early modern theatre, thus rendering it manageable. For example, stripped-down manuscript parts that excised all unnecessary information (including the other parts, save for sparse line cues) were used in conjunction with what were called ‘plots’ – sheets of paper containing scene-by-scene accounts of entrances and exits, casting, and

sound and music cues. This external scaffolding worked by assuming both the particular three-dimensional organization of the physical theatrical space (e.g. the door arrangements on the early modern stage), and certain conventions of movement that were operative in theatre at the time (meaning that the door through which an actor is to enter or exit is hardly ever specified in the aforementioned plots). Finally, various guild-like social structures and protocols supported the development of apprentice actors, enabling them to perform progressively more complex roles.

How does the distributed cognition paradigm bear on our present, archaeological concerns? Given that the environmental factors highlighted by the distributed perspective include items of material culture, it is tempting to see a symbiotic connection between that cognitive-scientific approach and archaeology, especially in contexts where archaeologists seek to draw conclusions about the human mind from material-cultural evidence. The background picture at work in the distributed perspective is of 'our distinctive universal human nature, insofar as it exists at all, [as] a nature of biologically determined openness to deep, learning- and development-mediated, change' (Wheeler and Clark 2008, 3572) and thus, given a technologically saturated environment, of human organisms as what Clark (2003) calls *natural born cyborgs*, creatures who are naturally evolved, both now and in the ancient past, to seek out intimate couplings with the non-biological resources of material culture. Crucially, as should be clear already, such couplings extend to the achievement of psychological feats such as, for example, memory and reasoning. Moreover, in a process that Clark has dubbed *cognitive niche construction* (e.g. Clark 2008; see also Wheeler and Clark 2008), which is nicely illustrated by the foregoing example of distributed memory in the early modern theatre, human beings design and build external structures that, often in combination with culturally transmitted practices, transform problem spaces in ways that typically promote, but sometimes (when things go wrong) obstruct, thinking and reasoning.

It is tempting to propose that when the problem of ancient minds is approached from the distributed perspective, that problem will lose much of its bite. After all, the idea that human cognition is shaped fundamentally by organic-technological hybridization might reasonably indicate that we ought to be able to say quite a lot about ancient minds from the archaeological evidence provided by the technology and material culture with which those minds were, by hypothesis, intimately coupled. Indeed, within recent archaeological theorizing, it might seem to be the explicitly distributed dimension of Malafouris's increasingly influential *material engagement theory* that enables him to use material culture as a productive bridge between experiential, conceptual, and social aspects of the self (Malafouris 2013). And one can feel the same putatively beneficial mutuality bubbling away just beneath the surface in the argument by Hodder in Chapter 1 regarding, specifically, cognition in the Neolithic period that, given the notion of a contextually distributed and plastic mind, it is reasonable to expect that the Neolithic, with its panoply of new techniques and ways of life, would be associated with cognitive change.

In this chapter, I shall interrogate the apparent symbiosis between the archaeology of cognition and the distributed perspective. I shall argue that, once one accepts that cognition is distributed, settling the nature of ancient thought, and, more particularly, determining whether cognitive change has occurred, on the basis of material-cultural evidence, is a rather more complicated business than one might have imagined. My case study will be the

alleged emergence of higher levels of cognition and consciousness in the Neolithic, as indicated by the material-cultural evidence at Çatalhöyük.

Before we start our investigation proper, two clarifications are in order. Firstly, In the spirit of cross-disciplinary fertilization, this chapter approaches the question of changes in consciousness and cognition at Çatalhöyük from the perspective of the philosophy of cognitive science. As a consequence, the focus of discussion will be on the underlying assumptions and the general structure of certain considerations and inferences, rather than on the detailed interpretation of specific pieces of material evidence. That said, and as will become clear, the arguments of this chapter, if correct, will have repercussions for how we interpret the material evidence, and, in particular, for the conclusions that we might draw, from the archaeological data at Çatalhöyük, regarding changes in the Neolithic mind. Put another way, the goal of this investigation is not merely to draw some theoretical conclusions about the relationship between distributed cognition and archaeology, but to extract some lessons about what we are entitled to say about changes in consciousness and cognition at Çatalhöyük, on the basis of the available material-cultural evidence, once we allow our hypotheses about ancient minds to be shaped by certain general accounts of the basic character of psychological phenomena that are operative in cognitive science.

Secondly, although the intention here is to argue that which theory of mind an archaeologist selects, from the available options in cognitive science, will have implications for what precisely can and cannot be said about ancient consciousness and cognition on the basis of the material record, that argument, even if successful, in no way blocks off the other direction of potential cross-disciplinary travel, that is, a situation in which what we learn from the material record leads us to say something different about our scientific theories of mind. We might discover, for example, that certain elements that we would expect to observe in ancient material culture, if a particular scientific theory of mind applies to that context, simply can't be found, thereby casting doubt on the applicability of that theory, at least as it stands. But that direction of travel is not what is at issue here. With those two clarifications in place, it is time for the real work of this chapter to begin.

2. Material Symbols and Cognitive Change

It is, of course, an immensely plausible thought that historical changes in habitation, agriculture, technology or trade indicate the introduction of new techniques and ways of life that will be associated with cognitive change. But that's a pretty 'big-picture' thought and things are not so obvious when we turn our attention to more detailed hypotheses about which particular changes or innovations in which particular practices drove, were driven by, or were at least correlated with, which particular cognitive changes. Here we will be taking our cues from one such detailed hypothesis.

Small geometric clay objects (spheres, discs, cones) are common finds in all occupational levels at Çatalhöyük. So what, precisely, are these objects? One prominent proposal is that they are 'tokens', or what (drawing on terminology from the distributed cognition literature) I shall call 'material symbols'. In other words, by functioning within a structured system of representational relationships, such physical objects become material resources for symbolic information storage. Bennison-Chapman, who, as we shall see, proceeds to

question this interpretation of the objects concerned, characterizes ‘tokens’ (material symbols) as: ‘small tools acting as mnemonic aids, used to hold and transmit information. They are utilised within the sphere of administration, to store and communicate information.’ (Bennison-Chapman, this volume, Chapter 5). So, if the small geometric clay objects recovered at Çatalhöyük are indeed material symbols, they should be interpretable as the physical realizers of recombinable symbols within a standardized system in which simple geometric shapes in different sizes are used to represent units of, for example, animals, crops, processed foods and raw materials.

So far, so good, but how does the claim that there are material symbols at Çatalhöyük bear on the question of cognitive change? According to thinkers such as Schmandt-Besserat (1992), Renfrew (1998, 2012) and Watkins (2010), the appearance of material symbols in the archaeological record is evidence of a shift in cognitive capacities, since only a more ‘advanced’ mind has the ‘higher’ cognitive abilities that are required to invent and operate a symbolic system, where the candidates for such ‘higher’ abilities include objectification, abstraction (taking a particular shape to represent, say, an animal) and metrication. With Neolithic village life in general (not just at Çatalhöyük) identified by the aforementioned thinkers as a source for such evidence, the precise historical drivers for the transition in question are a matter of debate, with some (e.g. Schmandt-Besserat, Renfrew) tending to favour something like a shift from a mobile hunter-gatherer lifestyle to one of sedentary agriculture, and thus from a form of community in which the counting of resources was not necessary to one in which it was, and others (e.g. Watkins) favouring factors such as the possibilities for expressing and storing more complex concepts opened up by the new built environment and the cognitive challenges posed by the social fact that the members of the community were now living alongside one another on a more permanent basis. This debate over the historical drivers of cognitive change (and it is important to note that the views I have just canvassed do not exhaust the available options), provides part of the backdrop to the archaeological research project explored in this volume, in which Çatalhöyük, with its large amount of data, is used as a laboratory for testing hypotheses about the causes of cognitive change (Hodder, this volume, Chapter 1). More specifically, one line of investigation is to agree that material symbols exist at the site and then to interrogate the archaeological data with the aim of revealing the cause of their emergence and thus of the associated cognitive change (Chapter 1). But it is the shared, motivating idea that sits behind the search for the specific historical drivers that concerns us here. That idea is that the deployment of material symbols requires certain sophisticated forms of cognition and consciousness, so the appearance of such symbols in material culture marks the emergence of new psychological capacities.

To help us explore this theoretical territory, and concentrating specifically on Çatalhöyük, we can express the central line of reasoning just traversed in the form of an explicit argument.

Premise 1: The appearance of material symbols in the archaeological record at a site constitutes evidence of cognitive change, in that period, at that site. More specifically, it constitutes evidence of a shift to certain sophisticated forms of cognition and consciousness.

Premise 2: The small geometric clay objects found at the Neolithic site of Çatalhöyük are material symbols.

Conclusion: There is evidence of cognitive change at the Neolithic site of Çatalhöyük, and, more specifically, of a shift to certain sophisticated forms of cognition and consciousness.

Now, if someone wanted to reject this argument, and thus the claim that there was a shift to sophisticated levels of cognition and consciousness at Çatalhöyük evidenced by the discovery of material symbols at the site (there might, of course, be an alternative basis for the conclusion that such a shift has taken place), one strategy would be to reject the second premise, that is, to argue that the small geometric clay objects found at Çatalhöyük are not material symbols. This is precisely what Bennison-Chapman (this volume) does. First she identifies a suite of conditions that would plausibly need to be met, if the clay objects in question really did function as material symbols in the manner required. Among these conditions are things like being crafted into a range of standardized shapes and sizes consistent with the range of goods present in the Çatalhöyük economy, being used in groups, and being retained for later information retrieval. Then she presents detailed archaeological evidence from the site which strongly suggests that the conditions in question are not met. For example, there is no evidence of development in the range or homogeneity of form of the clay objects. In addition, the objects in question are usually recovered alone and often within disposal contexts. Finally, she presents considerations in favour of certain alternative hypotheses regarding the function of the objects, such as that they were gaming pieces or that they were simple counting (as opposed to accounting or recording) tools, used in a one-to-one correspondence with individual items in particular counting events (and thus not as part of a symbolic system).

As fascinating and as important as this dispute is, it does not matter, for our purposes here, whether Bennison-Chapman is right that the small geometric clay objects found at Çatalhöyük are not material symbols. What's more significant is that that conclusion, which constitutes the denial of premise 2 in our highlighted argument, leaves premise 1 of that argument intact. Put another way, Bennison-Chapman is, as far as I can tell, happy to believe, along with her opponents, that the appearance of material symbols in the archaeological record at a site constitutes evidence of cognitive change – and more specifically of a shift to certain sophisticated forms of cognition and consciousness – in that period, at that site. Thus, for Bennison-Chapman and her opponents, *if* the geometric clay objects found at Çatalhöyük are material symbols, that would be good evidence of the emergence of certain advanced cognitive capacities there; it's just that, for Bennison-Chapman, those objects aren't material symbols, whereas, for her opponents, they are. But now what about the shared premise itself? Is *that* correct? If it isn't, then the link between material symbols and sophisticated levels of consciousness and cognition is broken, and the recovery of material symbols at some site would not be evidence of cognitive change. Under these circumstances, even if Bennison-Chapman is wrong, and the small geometric clay objects recovered at Çatalhöyük are indeed material symbols, that would not be evidence of the proposed emergence of sophisticated levels of consciousness and cognition in the Neolithic mind.

The remainder of this chapter will focus on premise 1 of the target argument, on, that is, the claim that the appearance of material symbols in the archaeological record at a site constitutes evidence of cognitive change, in that period, at that site, and more specifically of a shift to certain sophisticated forms of cognition and consciousness. Facing up to the problem of ancient minds, what we really want to know is what went on in the minds of our ancient ancestors when they interacted with external symbols systems, and whether the introduction of such systems heralded a radical shift in the fundamental nature of their cognitive resources. In an attempt to meet these challenges, we can pursue a strategy that is, in truth, a close neighbour of the neuroarchaeologist's appeal to what we know about the psychological states and processes that are present when contemporary human beings perform certain kinds of task (see above). That is, we can find out what contemporary cognitive science tells us about what goes on in people's minds when they interact with modern external symbol systems. More specifically, given that the distributed perspective looks like it's the archaeologist's friend, we can find out what that particular perspective in contemporary cognitive science tells us about the psychological states and processes in play. Then we can extrapolate to the case of ancient material symbols and ancient minds. So, it is by seeing if/how the distributed perspective provides a theoretical bridge from ancient material symbols to ancient cognition – a connection which is at the centre of the highlighted research on Çatalhöyük – that the alleged symbiosis between the archaeological investigation of the ancient mind and distributed cognition may be assessed.

3. Distributed Cognition 1: Embedded Minds

In a series of compelling cognitive-scientific treatments that combine philosophical reflection with empirical modelling studies, Bechtel (1994, 1996; see also Bechtel and Abrahamsen 1991) develops and defends the view that certain 'advanced' cognitive achievements, such as mathematical reasoning, natural language processing and natural deduction, are the result of sensorimotor-mediated embodied interactions between in-the-head connectionist networks and external symbol systems. In cognitive science, the term 'connectionist network' picks out a class of systems in which a (typically) large number of interconnected units process information in parallel. In as much as the brain too is made up of a large number of interconnected units (neurons) that process information in parallel, connectionist networks are 'neurally inspired', although usually at a massive level of abstraction. At the heart of connectionist theorizing is the concept of a distributed representation – a pattern of activation spread out across a group of processing units (analogous to a pattern of neural activity in a brain). One feature of connectionist networks that has made them popular with many cognitive theorists is that, as a by-product of their basic processing architecture and form of information storage, properties which endow them with a powerful line in statistical pattern completion, these systems 'naturally' demonstrate a range of intelligence-related capabilities that plausibly underlie the distinctive psychological profile of biological thinkers. Such capacities include flexible generalization from existing data, default reasoning, and the graceful degradation of performance in the face of restricted damage or noisy/inaccurate input information. What's striking about Bechtel's explanation of our sophisticated cognitive achievements in (roughly) mathematics, language and formal logic is that the genuinely psychological contribution at work is exhausted by these sorts of biologically realistic capabilities. The rest is a matter of embodied interaction with, and environmental scaffolding by, external material symbols.

But why, exactly, is this striking? To answer that question, and to understand more precisely, from a distributed perspective, what's going on in Bechtel's explanation, we need to make contact with one of the most famous quarrels in cognitive-scientific history.

In contrast to connectionism, the classical form of cognitive science denies that the abstract structure of the brain is a good model for the nature of mind. Rather, it demands, we should pay attention to the abstract structure of human language. Such language (on one popular account anyway) is at root a finite storehouse of atomic symbols (words) which are combined into complex expressions (phrases, sentences, and so on) according to certain formal-syntactic rules (grammar). The meaning of some complex expression is a function of the meaning of each atomic symbol that figures as a constituent in that expression, plus the syntactic structure of the expression (as determined by the rules of the grammar). In short, human language features a combinatorial (equivalently, compositional) syntax and semantics. And, for the classical cognitive scientist, so it goes for our inner psychology. That too is based on a finite storehouse of atomic symbols (concepts) that are combined into complex expressions (thoughts) according to a set of syntactic rules. In short, thinking, like language, features a combinatorial syntax and semantics. Thus Fodor famously speaks of our inner psychological system as a *language of thought* (Fodor 1975).

One much-discussed argument for the classical view (and thus against connectionism) hails from Fodor and Pylyshyn (1988) who claim that connectionist theorizing about the mind is, at best, no more than a good explanation of how classical states and processes may be implemented in neural systems. In brief, Fodor and Pylyshyn argue as follows. It's an empirical observation that thought is systematic. In other words, the ability to have some thoughts (e.g. that Elsie loves Murray) is intrinsically connected to the ability to have certain other thoughts (e.g. that Murray loves Elsie). If we adopt a classical vision of mind, the systematicity of thought is straightforwardly explained by the combinatorial syntax and semantics of the cognitive representational system. The intrinsic connectedness of the different thoughts in question results from the fact that the processing architecture contains a set of atomic symbols alongside certain syntactic rules for recombining those symbols into different molecular expressions. Now, Fodor and Pylyshyn argue that although there is a sense in which connectionist networks instantiate structured states (e.g. distributed connectionist representations have active units as parts), specifically combinatorial structure is not an essential or a fundamental property of those states. This ultimately renders connectionist networks inherently incapable of explaining the systematicity of thought, and thus of explaining thinking. What such systems might do, however, is explain how a classical computational architecture may be implemented in an organic brain.

Bechtel explicitly develops his account in opposition to Fodor and Pylyshyn's deflationary conclusions regarding connectionism. That said, he agrees with Fodor and Pylyshyn on two key points, firstly that where systematicity is present, it is to be explained by combinatorially structured representations, and secondly that connectionist networks fail to instantiate combinatorial structure as an essential property of their internal organization. He does not need to endorse Fodor and Pylyshyn's claim that all thought is systematic, however. For his purposes, all that is required is that some cognitive activities (e.g. linguistic behaviour, natural deduction, mathematical reasoning) exhibit systematicity. Bechtel's distinctive (anti-

Fodor-and-Pylyshyn) move is to locate the necessary combinatorial structure in systems of symbolic representations that remain *external* to the connectionist network itself (e.g. in written or spoken language, or in mathematical/logical notations). Given the idea that our inner psychology should be conceived in connectionist terms, this is tantamount to saying that the necessary combinatorial structure resides not in our internal processing engine, but rather in the environment. For this solution to work, it must be possible for the natural sensitivity to statistical patterns that we find in connectionist networks generally to be deployed in such a way that some of those networks, when in interaction with specific external symbol systems, may come to respect the constraints of a combinatorial syntax, even though their own inner representations are not so structured. Bechtel's studies suggest that this may be achieved by exploiting factors such as the capacity of connectionist networks to recognize and generalize from patterns in bodies of training data (e.g. large numbers of correct derivations in sentential arguments), plus the temporal constraints that characterize real embodied engagements with stretches of external symbol structures (e.g. different parts of the input will be available to the network at different times, due to the restrictions imposed by temporal processing windows). The conclusion is that 'by dividing the labor between external symbols which must conform to syntactical principles and a cognitive system which is sensitive to those constraints without itself employing syntactically structured representations, one can perhaps explain the systematicity... of cognitive performance' (Bechtel 1994, 438).

As defined earlier, cognition is distributed when it is, in some way, spread out over the brain, the non-neural body and (in many paradigm cases) an environment consisting of objects, tools, other artefacts, texts, individuals and/or social/institutional structures. Bechtel's model of an inner connectionist network in embodied interaction with material symbols, where the external symbol system explains the systematicity of cognitive performance, fits the bill. Moreover, it's an example of the distributed perspective at work that speaks directly to the issue of what goes on in people's minds when they interact with external symbol systems. But now we need to be more specific about the precise form of distributed cognition that's on the table, because it turns out that, despite the shared, paradigm-defining emphasis on bodily engagement and environment-involving processing, the term 'distributed cognition' is actually an umbrella concept that encompasses a number of distinct theoretical views. Two modulations of the core idea will be relevant here, modulations which are generated by different conceptions of what is involved in cognition *spreading out* over brain, body and world.

According to the hypothesis of *embedded* cognition, the distinctive adaptive richness and/or flexibility of intelligent thought and action is regularly, and perhaps sometimes necessarily, causally dependent on the bodily exploitation of certain environmental props or scaffolds. (For philosophical elucidations and explorations of this idea, see e.g. Clark 1997, Wheeler 2005), What's important about this approach, for present purposes, is that although the embedded theorist seeks to register the routinely performance-boosting, often transformative, and sometimes necessary causal contributions made by environmental elements to many cognitive outcomes (witness the role of Bechtel's external symbols in explaining systematicity), she continues to hold that the actual thinking going on in such cases remains brain-bound. (There is a less common, more radical iteration of the view according to which cognition is distributed through the brain *and* the non-neural body,

although not the environment. I shall ignore this option.) It's the embedded version of distributed cognition that Bechtel seems to prefer. Thus his claim that the 'property of systematicity, and the compositional syntax and semantics that underlie that property, might best be attributed to natural languages themselves but not to the mental mechanisms involved in language use' (Bechtel 1994, 436) is plausibly explained by the fact that he takes the only genuinely cognitive elements in the nexus of neural, bodily and environmental factors to be inside the head: systematicity and combinatorial structure are features of the external symbol systems, but not the mental mechanisms involved, because the only mental mechanisms involved are the inner connectionist networks that do not themselves exhibit systematicity or combinatorial structure.

Re-entering the archaeological context, if the embedded interpretation of Bechtel's model is correct, what we confront is an infuriating resuscitation of the problem of ancient minds. To see why, notice that, on that interpretation, the cognitive mechanisms implicated in navigating and exploiting the sorts of external material symbol system under consideration (material realizations of symbolic logical notation, say) do not involve sophisticated cognitive capacities, over and above those that were already operative in the non-symbolic context. So, by extrapolation, the same natural sensitivity to statistical patterns that we find in connectionist networks generally, a sensitivity which underlies capacities for generalization, default reasoning and graceful degradation that surely will have been central to all kinds of adaptive thought and experience in prior hunter-gatherer communities, may continue to characterize cognition in the new Neolithic environment of (we are currently assuming) material symbols, simply by becoming targeted on patterns in the external symbolic systems in question. So, if, from the perspective of the embedded version of distributed cognition, we ask ourselves the question 'Does the presence of an external material symbol system require, or at least strongly indicate, cognitive change, in the sense of the installation of a revolutionary package of more sophisticated psychological capacities?', our answer should be 'no'. This is because, on the basis of Bechtel's model, the increased sophistication of the observed psychological *performance* may have been purchased using the same old connectionist currency of statistical pattern completion that was already operative, but which is now newly allied with some powerful external scaffolding (the material symbol system) that allowed the Neolithic mind, including that mind as realized at Çatalhöyük, to buy more for its psychological money.

These thoughts find additional support in some reflections by Clark (1997) on what happens to the brain with the advent of language. Clark's proposal – which we can interpret for the moment as an example of embedded theorizing – is that language should be thought of as 'an external resource that complements but does not profoundly alter the brain's own basic modes of representation and computation' (198). In other words, the biological brain has certain generic forms of inner state and mechanism ('the brain's own basic modes of representation and computation') that, from both an evolutionary and a developmental perspective, precede linguistic competence. When language comes onto the cognitive scene, it heralds not a transformation in those generic types of inner resource, but rather an external augmentation of them. Indeed, buying into a broadly connectionist approach to cognitive science, Clark takes the human brain to be essentially a device for pattern-association, pattern-completion and pattern-manipulation. So Clark's recognizably Bechtelian claim is that our language-involving behaviour is to be explained by an all-

conquering partnership between, on the one hand, a pattern-sensitive brain and, on the other, an external storehouse of rich symbolic structures.

To bring Clark's position into proper view, it is worth pausing to compare it with a related account that, tentatively, Clark himself attributes to Dennett (Dennett 1991; Clark 1997, 197). Dennett argues that our innate neural hardware may differ from that of our non-linguistic evolutionary near-neighbours (such as chimpanzees) in only relatively minor ways. Nevertheless, it is precisely these relatively minor hardware differences that constitute the evolutionary source of the human ability to create, learn and use public language. According to Clark, this part of Dennett's story is correct: there is no mandate to attribute human beings with the kind of innate language processing mechanism whose design would mean that our brains, compared with those of our evolutionary near-neighbours, contain a fundamentally different kind of neural device. However, Dennett's further proposal, as Clark explains it, is that developmental exposure to a linguistic environment results in a subtle reprogramming of the computational resources of the human brain, such that our innate pattern-completing neural architecture comes to simulate a kind of logic-like serial processing device. Clark, by contrast, resists the idea of any extensive ontogenetic reprogramming phase driven by language. Thus, we are told, developmental exposure to and use of language brings about no significant reorganization of the brain's basic processing architecture. If this is true for our linguistic capabilities, the same points can surely be made regarding the psychological capacities required to create, learn and use the material symbol systems that pre-date writing.

The foregoing considerations cast doubt on the inference from the appearance of material symbols in the archaeological record to the occurrence of significant cognitive change. They also indicate that distributed cognition may not be quite the friend that the archaeologist interested in ancient minds needs. After all, so far anyway, our foray into distributed cognition suggests that, with the explanatory weight spread over brain, body and world, plus, crucially, the intimate causal couplings between these different but complementary elements, there is no secure inferential bridge from what we recover in the form of material culture to changes in the nature of cognition. Old brains don't necessarily need to learn substantially new cognitive tricks, in order to deliver more sophisticated psychological achievements following innovations in material culture.

4. Distributed Cognition II: Extended Minds

Perhaps we are not being radical enough in our adoption of a distributed perspective. In spite of all its exciting talk of embodied interaction and environmental scaffolding, the fact remains that the embedded version of distributed cognition shares a central traditional assumption with the neuro-centric orthodoxy, namely that psychological states and processes are instantiated in physical machinery that is *always* inside the skull and skin. Thus certain external elements act as non-cognitive factors that support and augment the wholly internal cognitive states and processes. Looking at this aspect of the embedded view from a different angle, and borrowing a way of putting things from Adams and Aizawa (2008), a key feature of embedded cognition is that the dependence of cognition on external elements is causal – or, to bring out what really matters, *merely* causal – rather than *constitutive*. If the dependence in question were constitutive in character, then the

external elements of interest would not be non-cognitive causal scaffolds for cognition; they would count as genuine *parts of* the cognitive process or architecture.

This is the view endorsed by the second version of the distributed perspective to be canvassed here. Thus, according to the hypothesis of *extended* cognition, the physical machinery of mind sometimes extends beyond the skull and skin (see Clark and Chalmers 1998 and Clark 2008 for canonical treatments, and for a more recent collection that contains criticisms, defences and developments of the view, see Menary 2010). More precisely, the advocate of extended cognition holds that there are actual (in this world) cases of intelligent thought and action, in which the material vehicles that realize the thinking and thoughts concerned are spatially distributed over brain, body and world, in such a way that certain external elements are rightly accorded fundamentally the same cognitive status as would ordinarily be accorded to a subset of your neurons. If this view is right, then, under certain circumstances, your phone-number-storing mobile device literally counts as part of your mind, in the sense that it's part of your mnemonic machinery.

Here is not the place to develop and defend a detailed account of when some external element, such as a material symbol, qualifies as a constituent part of one's psychological machinery, that is, in the relevant sense, as part of one's extended mind. As we might expect, there are several proposals for delivering this result and a sometimes bad-tempered debate surrounds them. (Menary 2010 is a good place to make contact with the debate. For my own favoured way of arguing for extended cognition, see e.g. Wheeler 2010a, 2010b, 2011.) However, it's important to note that the transition to approaching things in terms of extended, rather than embedded, cognition does not necessarily require any associated change in the underlying causal structure of the distributed system under consideration. Rather, it may be based on an acknowledgment that some of the causal structures present in that system meet the conditions for cognitive status. To see how this might work, consider, once again, the Bechtelian connectionist-network-plus-symbol-system architecture described above, which so far, and in line with Bechtel's own approach, has been treated as an instance of embedded cognition. For the sake of argument, let's simply assume that we have been convinced by various considerations that where one finds a suitably organized material system in which atomic symbols are automatically combined and manipulated, according to the principles of a compositional syntax and semantics, so as to meet Fodor and Pylyshyn's systematicity requirement, one finds a cognitive system. (Something like this view is endorsed by Newell and Simon's physical symbol system hypothesis, one of the canonical statements of classical cognitive science; Newell and Simon 1976.) Under these circumstances, one might be moved to claim that Bechtel's distributed architecture, comprising an inner connectionist network and external symbols coupled together via embodied sensorimotor interactions, itself qualifies as just such an automatic, material, compositional symbol system, and thus as a cognitive system. In other words, in this distributed architecture, the genuinely cognitive machinery includes not only the connectionist network, but also the embodied interface and the external material symbols themselves. Conceived this way, a Bechtelian architecture for doing mathematics, natural language processing or natural deduction is a case of extended cognition.

If the extended interpretation of Bechtel's model is correct, things certainly look a little better for the archaeologist interested in ancient minds. After all, whether we are working

from within neuro-centrism or from within an embedded distributed perspective, material culture is condemned to a life outside of cognition proper, and so the objects and artefacts studied by archaeology are (roughly) things that ancient minds made and/or used. The cognitive states and processes concerned are not themselves simply on show in those things, although certain inferences about the nature of those states and processes might be ventured. (Enter the problem of ancient minds.) However, with the cognitive architecture concerned now conceptualized as *including* the external material symbols, past ways of thought are not just *expressed in* material culture but are often partly *constituted by* material culture. So the archaeologist gets to study past minds in a rather more direct manner. Indeed, some of the things studied by archaeology are literally parts of (no longer functioning) ancient minds (Wheeler 2010b). So, if, from the vantage point afforded by the extended interpretation of Bechtel's architecture, we ask ourselves the question 'Does the presence of an external material symbol system require, or at least strongly indicate, cognitive change, in the sense of the installation of a revolutionary package of more sophisticated psychological capacities?', then our answer should be 'yes'. This is because the appearance of material symbols to which certain internal mechanisms are coupled via embodied sensorimotor interactions is itself (it is not merely evidence of) cognitive change. It is the emergence of new and more sophisticated cognitive structures. This remains true even if (and we cannot be sure about this) the inner part of the extended ancient mind instantiates nothing more than the old connectionist currency of statistical pattern completion that was already operative in the pre-symbolic context. So, for fans of the extended cognition hypothesis, if the small geometric clay objects at Çatalhöyük are in fact material symbols, then the Neolithic mind at that site did undergo significant cognitive change.

It might seem as if our extended-mind-based argument for the conclusion that cognitive change has occurred has the unhappy effect of rendering that conclusion empty or trivial. The complaint here would go as follows: if material symbols routinely count as cognitive elements, and if the notion of the 'symbolic' applies generally to many elements in material culture, then a very large number of the changes in material culture will count straightforwardly as cognitive changes. So the conclusion that cognitive change has occurred says nothing – or at least nothing interesting or illuminating. Fortunately, the specific extended-mind-based argument outlined earlier does not fall prey to this particular objection, because that argument is developed in the context of (i) material elements that are symbolic in a restricted and demanding sense, in that they function as recombinable constituents within standardized systems of representational structures and relationships, and (ii) a related, and equally restricted and demanding, account of 'cognition', according to which certain outcomes in human performance are underpinned by atomic symbols that are exploited and manipulated according to the principles of a compositional syntax and semantics, thus enabling the systematicity condition for thought to be met in those domains. As Bennison-Chapman nicely shows (see above), not all material-cultural elements will be material symbols *in the sense of (i)*. Building on this, we can now add, given (ii), that the kinds of exploitation and manipulation to which the (by our current criterion) non-symbolic elements will be subjected will be of the *wrong kind* to count as cognitive. That should be enough to alleviate the trivialization concern, because a large number of the changes that take place in material culture will not count as cognitive changes. In other

words, if one concludes, on the basis of the extended-mind-based argument offered earlier, that cognitive change has occurred, one is saying something with genuine content.

5. Conclusions

In this chapter, I have used the example of the alleged emergence of higher levels of cognition and consciousness in the Neolithic, as indicated by the material-cultural evidence at Çatalhöyük, to explore the relationship that exists between the archaeology of the ancient mind and distributed cognitive science. I have argued that the conclusions that should be drawn regarding the occurrence of cognitive change, based on the material-cultural evidence, will be different (indeed, diametrically opposed) depending on which version of distributed cognition is embraced. Adopting the embedded version, one cannot know whether there has been cognitive change on the basis of the appearance of material symbols in the archaeological record, since such change may not be required to explain our successful deployment of those symbols. Adopting the extended version, one may be able to conclude that cognitive change has occurred, since the material symbols themselves may have cognitive status, but that conclusion would still leave one unsure about the character of the inner contribution to the overall cognitive process. So, although there may be good reason to think that the distributed perspective is the archaeologist's friend, that enthusiasm needs to be matched by an understanding of the challenges, as well as the advantages, of such an alliance.

Acknowledgments

Some passages in section 3 of this paper were adapted from passages in (Wheeler 2004, 2005, 2015). For discussion of the ideas presented here, many thanks to the participants at the conference, *Consciousness and Creativity at the Dawn of Settled Life*, McDonald Institute for Archaeological Research, University of Cambridge, July 2017, and especially to Ian Hodder and John Sutton.

References

Adams, F. and Aizawa, K. (2008), *The Bounds of Cognition*. Blackwell: Malden, MA and Oxford.

Bechtel, W. (1994), "Natural Deduction in Connectionist Systems", *Synthese* vol. 101, pp.433-463.

Bechtel, W. (1996), "What Knowledge must be in the Head in order to Acquire Language", in *Communicating Meaning: the Evolution and Development of Language*, eds. B. Velichkovsky and D. M. Rumbaugh, Lawrence Erlbaum Associates: Hillsdale, NJ.

Bechtel, W. and Abrahamsen, A. (1991), *Connectionism and the Mind: an Introduction to Parallel Processing in Networks*, Basil Blackwell: Oxford.

Bennison-Chapman, L.E. (this volume), "Conscious Tokens?".

- Clark, A. (1997), *Being There: Putting Brain, Body, And World Together Again*, MIT Press: Cambridge, Mass.
- Clark, A. (2003), *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*, OUP: New York.
- Clark, A. (2008), *Supersizing the Mind: Embodiment, Action and Cognitive Extension*, OUP: Oxford.
- Clark, A. and Chalmers, D. (1998), "The Extended Mind", *Analysis*, Vol. 58, no. 1, pp.7-19.
- Dennett, D.C. (1991), *Consciousness Explained*, Little, Brown & Co.: Boston.
- Fodor, J.A. (1975), *The Language of Thought*, Harvard University Press: Cambridge, Massachusetts.
- Fodor, J.A. and Pylyshyn, Z. (1988), "Connectionism and Cognitive Architecture: a Critical Analysis", *Cognition*, Vol. 28, pp.3-71.
- Hutchins, E. (1995), *Cognition in the Wild*, MIT Press: Cambridge, Mass..
- Malafouris, L. (2013), *How Things Shape the Mind*. MIT Press: Cambridge MA.
- Menary, R. (ed.) (2010), *The Extended Mind*. MIT Press: Cambridge, Mass.
- Mithen, S. (1996), *The Prehistory of the Mind: A Search for the Origins of Art, Religion and Science*, Thames and Hudson: London.
- Newell, A. and Simon, H. A. (1976), "Computer Science as Empirical Inquiry: Symbols and Search", *Communications of the Association for Computing Machinery*, vol. 19, no. 3, pp.113-126.
- Putt, S.S., Wijekumar, S., Franciscus, R. G. and Spencer, J.P. (2017), "The Functional Brain Networks that Underlie Early Stone Age Tool Manufacture", *Nature Human Behaviour*, vol. 1, Article number: 0102, doi:10.1038/s41562-017-0102
- Renfrew, C. (1998), "Mind and Matter: Cognitive Archaeology and External Symbolic Storage" in *Cognition and Material Culture: The archaeology of symbolic storage*, eds. C. Renfrew & C. Scarre, McDonald Institute for Archaeological Research: Cambridge, pp. 1-6.
- Renfrew, C. (2012), "Towards a Cognitive Archaeology: Material Engagement and the Early Development of Society", in *Archaeological Theory Today*, ed. I. Hodder. Polity Press: Cambridge and Malden, Mass., pp. 124-45.
- Schmandt-Besserat, D. (1992), *Before Writing (Vol. 1): From Counting to Cuneiform*, University of Texas Press: Austin.

Tribble, E. B. (2005), "Distributing Cognition in the Globe", *Shakespeare Quarterly*, vol. 56, no. 2, pp.135-155.

Tribble, E. B. (2011), *Cognition in the Globe: Attention and Memory in Shakespeare's Theatre*, Palgrave Macmillan: New York.

Watkins, T. (2010), "New Light on Neolithic Revolution in South-West Asia", *Antiquity*, vol. 84, no. 325, pp. 621-34.

Wheeler, M. (2004), "Is Language the Ultimate Artefact?", *Language Sciences*, vol. 26, no. 6, pp. 693-715.

Wheeler, M. (2005), *Reconstructing The Cognitive World: The Next Step*. MIT Press: Cambridge, Mass..

Wheeler, M. (2010a), "In Defense of Extended Functionalism", in (Menary 2010), pp.245-270.

Wheeler, M., (2010b), "Minds, Things and Materiality", in *The Cognitive Life of Things: Recasting the Boundaries of the Mind*, eds. C. Renfrew and L. Malafouris, McDonald Institute for Archaeological Research: Cambridge, pp.29-37.

Wheeler, M. (2011), "Embodied Cognition and the Extended Mind", *The Continuum Companion to Philosophy of Mind*, ed. J. Garvey, Continuum: London, pp.220-238.

Wheeler, M. and Clark, A. (2008), "Culture, Embodiment and Genes: Unravelling the Triple Helix", *Philosophical Transactions of the Royal Society*, vol. 363, pp.3563-75.