Phenom Cogn Sci (2010) 9:465–481
DOI 10.1007/s11097-010-9174-y

Minds: extended or scaffolded?

Kim Sterelny

Published online: 5 October 2010

Abstract This paper discusses two perspectives, each of which recognises the importance of environmental resources in enhancing and amplifying our cognitive capacity. One is the Clark–Chalmers model, extended further by Clark and others. The other derives from niche construction models of evolution, models which emphasise the role of active agency in enhancing the adaptive fit between agent and world. In the human case, much niche construction is epistemic: making cognitive tools and assembling other informational resources that support and scaffold intelligent action. I shall argue that extended mind cases are limiting cases of environmental scaffolding, and while the extended mind picture is not false, the niche construction model is a more helpful framework for understanding human action.

Keywords Extended mind · Externalism · Niche construction · Social learning

The world made mental

In the last 15 years, Andy Clark, David Chalmers, Dan Dennett, Robert Wilson, Richard Menary, John Sutton and others have built a case for a new way of thinking about the mind and its boundaries. They have noticed how intimately human cognition—in particular, cognition that leads to competent action in the world—depends on environmental resources. We deal competently with problems that require quantitative assessment of our circumstances in part because we can measure our world with instruments; in part because we have developed public representational systems for representing quantitative information, systems with

K. Sterelny (✉)
Philosophy and Tempo and Mode, Australian National University, 0200 Acton, ACT, Australia
e-mail: kimbo@coombs.anu.edu.au

K. Sterelny
Philosophy, Victoria University of Wellington, Wellington, New Zealand
e-mail: Kim.Sterelny@vuw.ac.nz

efficient and user-friendly notation and in part because we have devices that transform quantitative representation rapidly and accurately. We respond competently to problems that require for their solution information we acquired long ago, in part because we can park information in books, diaries, post-it stickers and the like.

Clark and his allies are struck by such examples and in response have abandoned a model of the mind that sees it as an internal control system, enclosed in the human body, receiving data from human sensory systems and directing human action. Instead they see human minds as systems that extend beyond the body of the human organism, systems that include extra-somatic resources: environmental fuels for adaptive action. In general, the defenders of this "extended" view of the mind[1] do not suggest that a human's cognitive system includes all the environmental resources that support adaptive decision making. Rather, they suggest that human cognitive systems include those resources that are importantly, robustly, reliably or persistently supportive of decision making.

This paper explores the extended mind model by comparing it to an alternative that equally accepts the centrality of environmental resources to human intelligence. Many animals intervene in their environment, shaping it in ways that improve the adaptive fit between the agent and its world; such animals in part adapt to their niche; in part they construct their own niche. Many of the examples that motivate the extended mind picture can be seen as special cases of niche construction, cases in which human competences depend intimately on the environment being scaffolded to support adaptive decision making (Laland 2007; Laland et al. 2000; Odling-Smee et al. 2003). While I do not take this paper to refute the extended mind framework, I shall argue that the theory of niche construction offers a more insightful perspective on the role of environmental resources in human cognition, in part because it offers a more general perspective. I shall argue that the canonical extended mind cases are continuous with other cases, cases in which there is environmental support of cognition, but which are not plausibly treated as constituents of agents' minds. Moreover, the dependence of cognitive competence on extra-somatic resources turns out to be a special case of a more general phenomenon. Many human capacities, cognitive and non-cognitive alike, turn out to depend on the fact that humans engineer their environment to support their activities. The niche construction perspective focuses our attention on the common features of this whole range of cases; the extended mind model does not.

In the next section, I show the importance of environmental support for non-cognitive capacities and then introduce the niche construction perspective in more detail. Sections 4, 5 and 6 explore alternative forms of the role of external resources in cognition, making the case that the niche construction perspective is more insightful. I then briefly conclude.

[1] The beyond-the-body conception of the mind goes by many other names: "vehicle externalism"; "active externalism" and "locational externalism". I shall for simplicity use the "extended mind" terminology; I think it was the first used explicitly for this position and is the most widely recognised terminology. See, for example, (Clark 2007, 2008; Clark and Chambers 1998; Dennett 2000; Menary 2010; Sutton 2006; Sterelny 2010; Wilson 1994).

Extended digestion

The physiological demands on hominin jaws, teeth and guts have been transformed by cooking and more generally by food preparation and food targeting. Other great apes have to grind, shred and pulp what they eat with their teeth and jaws. They have to digest animal protein, fungi, fruits, nuts, seeds and other plant materials with their on-board equipment. As a consequence, they have powerful jaws, large mouths crammed with strong teeth and long guts. Even so, in order to meet their nutritional requirements, they spend a large amount of their waking day—4 to 5 h—just chewing (Wrangham 2009). In contrast, we variously transform our fuel by cooking, soaking, fermenting, grinding and pounding. We then reduce the mechanical demands of eating by chopping, slicing or breaking the prepared food into easily ingested, bite-sized pieces. These interventions vastly reduce the mechanical demands on efficient eating.

Cooking, soaking, fermentation and the like also change the chemistry of food in ways that enhance its food value. Richard Wrangham points out that raw-food enthusiasts find it very difficult to eat enough food, with enough variety, to maintain good health and biological function. In particular, men sometimes show much reduced sexual drive, and women stop ovulating. Raw foodists have this problem despite two huge advantages over ancient humans. They use blenders and similar technologies to liquefy and otherwise soften and/or disassemble food—and this in itself improves digestibility. Moreover, modern storage and transport makes a wide range of foods continually available. Even so, good health, weight and sexual function are hard to sustain (Wrangham 2009). We are obligatorily cooks. Moreover, we supplement cooking by pre-engineering our food sources. Over the last 10,000 years, selective breeding has improved the food value of domestic stocks (and reduced the toxins found in the tastier but more challenging wild types). Before domestication, hominin foragers improved the nutrition value of food before preparation by targeted foraging. Our distant ancestors selected heavily defended but high value foods (Gurven and Hill 2009; Hill and Kaplan 1999; Sterelny 2007). Not for us was the koala-like chore of internally fermenting leaves and similar roughage. Real men ferment externally.

There is plenty of debate about the onset of cooking in the hominin lineage. Wrangham, at one extreme, places it close to 2 Ma ago (Wrangham et al. 1999; Wrangham 2001). If he is right, the invention of cooking transformed human facial morphology, allowing teeth and jaw muscles to shrink, and helped to deliver the resources that made our large but expensive brain possible. Gut tissue is extensive too, so cooking, by reducing the demands on digestion, allowed hominins to reduce their investment in gut tissue, again making resources available for brain-building. But even if Wrangham's dates are too old, no one doubts that these technologies are ancient and important: There is clear evidence of hearths and controlled fire 400,000 years ago. Foley and Gamble do not place cooking as early as does Wrangham, but they think it might well be as early as 1 Ma ago, associated with the origins of *Homo heidelbergensis* and significant brain expansion in the human lineage (Foley and Gamble 2009).

Our digestion is, then, technologically supported in profound and pervasive ways. Our lives depend on the artefacts and techniques that make it physically possible for

us to ingest food and which enable us to extract more nutritional value from that food. We have engineered our gustatory niche; we have transformed both our food sources and the process of eating itself. Our under-powered jaws, short gut, small teeth and mouth fit our niche because we eat soft, rich and easily digested food. Our digestive system is *environmentally scaffolded*. But is my soup pot, my food processor and my fine collection of choppers part of my digestive system? As far as I know, no one has defended an extended stomach hypothesis, treating routine kitchen equipment as part of an agent's digestive system; indeed "extended stomach" and "extended digestion" on Google Scholar return some very strange hits.[2] As we have seen above, Andy Clark and his allies have suggested that external resources are sometimes genuine components of human cognitive systems, so perhaps silence on the extended stomach is an accident of intellectual history.

Identifying the role of environmental resources in cognition is the main target of this paper, but it will be useful to keep extended digestion in mind as a foil. It reminds us that environmental resources play an indispensible role in supporting many human capacities. Most obviously, tools and weapons amplify and transform many physical capacities. A tool as simple as a digging stick gives a hominin access to a whole new range of resources, for many plants have rich storage organs buried under hard, baked ground (O'Connell et al. 1999). The example of enhanced digestion is an especially important foil if one thinks (as is plausible) that there is no explanatory mileage in treating my soup pot as part of my digestive system, once its importance as a scaffold is recognised. Wrangham thinks that scaffolded digestion made us human: It was the crucial adaptive breakthrough that put our lineage on a path diverging from the other chimps (Wrangham 2009). Right or not, that idea shows we do not downplay the importance of external support in seeing digestion as enhanced by external resources. To recognise the importance of cooking, we need not suppose that digestion takes place on my stove as well as in my body.

A final point of target identification There is a third model for understanding how agent capacities are supported by environmental resources, namely Dawkins's conception of extended phenotypes (Dawkins 1982). Birds are great builders. Many species build strong, soft, weatherproof and often camouflaged nests of twigs, leaves, moss, spider web, fur and much else. For example, one Australian gerygone species constructs a nest that is suspended over water and which looks for all the world like a small clump of flood debris (Serventy 1982, pp 178–179). These hanging nests are part of the bird's phenotype. Like caddis fly houses, wasp nests, beaver dams and spider webs, they are adaptive complexes. They are as developmentally stable, as heritable and as predictable in their ecological effects as other traits. There are very likely "genes for" nest building traits in the nondeterministic sense of Dawkins 1982: genes that in typical genetic and developmental environments increase the probability of nest building behaviour relative to genetic alternatives at the same locus. The main point of Dawkins' extended phenotype conception is to emphasise the evolutionary and developmental

[2] Mark Rowlands has defended the view that there could be external digestion, but he has in mind much more recondite cases, in which processing that is normally managed by internal mechanisms depends instead on prosthetic devices (Rowlands 2009).

similarities between extended and standard phenotypic traits. Nest evolution presents an evolutionary problem of the same kind as other complex adaptations. To understand it, we need to decompose its evolutionary trajectory into a sequence of small changes, each of which, in the relevant environment, was an incremental improvement.

The precise relationship between extended phenotypes and constructed niches is contested, but at most a subset of niche construction cases fit Dawkins's model (Dawkins 2004; Laland 2004). In particular, neither my culinary examples nor Clark's cognitive examples fit the nest model. Most extended mind examples are examples of *adaptive plasticity*: of a novel yet adaptive phenotype. In his original presentation of the idea with David Chalmers, Clark developed the idea of the extended mind through the example of Otto, a man with memory problems. Otto coped with his inability to store new memories internally by keeping crucial information in a notebook, which he always kept with him. Otto's response does not depend on genes for notebook use, even in Dawkins's nondeterministic sense. Moreover, Otto's solution to the problem of memory loss was not built by cumulative improvement over the generations. I do not use Google as a search tool because my parents used it. Still less do I use it because my googling parents outcompeted Yahoo or Hotbot users. Contemporary wideware is not very heritable. Slide rules, pocket calculators, GPS devices, filofaxes and palm pilots appear in one generation and then disappear, sometimes within the same generation. Where are the filofaxes of yesteryear? Likewise, soup pots are a relatively new feature of the human landscape. Some technology is ancient, and there have been coevolutionary interactions between human phenotypes, human genomes and human technology (one may be the controlled use of fire), including perhaps informational technology. Our distant ancestors may have been marking trails and in other ways altering their environment to enhance their cognitive capacities. But we make and use many kitchen and cognitive tools adaptively, despite having no evolved disposition to use those very tools in those circumstances. The extended mind, as Clark, Wheeler, Wilson, Sutton and co-conceive of it, is not a special case of an extended phenotype. It is a special case of adaptive phenotypic plasticity.

Environmentally supported cognition

Over the last decade, I have been developing a model of the evolution and operation of human cognitive capacities, a model that places great weight on the role of the environment in supporting and amplifying our cognitive competences. The big picture idea depends on accepting nativist psychology's contention that our everyday cognitive skills depend on the mastery of large and often subtle bodies of information, while rejecting the idea that the core of these informational resources is innate, preinstalled by evolution. In my view, the adaptive problems confronting human agency are too variable at and across time for the nativist solution to be, in general, viable (Sterelny 2003, 2006).

The alternative I defend derives from the niche construction models cited earlier. Those models emphasise the active role of the agent in explaining the adaptive fit of

agent and environment. Over time, agents (or, more exactly, lineages of agents) do indeed adapt to their environments. But they also adapt their environment to them. In the classic examples of niche construction, these environmental modifications are directly utilitarian physical interventions. Animals construct nests, burrows and dams, thus protecting themselves from predators and from the violence of the world. Trees alter soils with their leaf litter, discouraging competitors. Earthworms modify the structure and moisture of the soils in which they live, easing the threat of desiccation. But epistemic action is a form of niche construction too. Agents alter the informational character of their environment in ways that make crucial features more salient. Thus, ants lay scent trails between nest and food source; sharks (supposedly) make feinting arracks on potential targets, safely testing their capacity to react; most simply, a hawk chooses a roost which maximises its view of its hunting territory. Moreover, niche construction often has cross-generation effects. Thus, John Odling-Smee and his colleagues think of it as a channel of ecological inheritance. Nests, burrows and beaver lodges are often reproductive resources, not just adult shelters. The animals that build these resources are engaging in downstream niche construction, reshaping the developmental environment of the next generation (Sterelny 2009).

Humans profoundly modify both the physical and the informational environment of the following generations. One way humans differ from the great apes is that intergenerational social learning profoundly shapes our minds and lives. I have recently argued that this social learning evolved around a core: the intergenerational transmission of ecological and technical expertise. This intergenerational flow began as a by-product of adult economic activity and then elaborated into something like an apprentice system. In the early stages of this evolutionary transition, children acquired parental skills as a side effect of the parental exercise of those skills. These early skills, I suspect, were important but not very complex. Likely examples are those skills involved in the production of the simplest flake-and-cobble stone technology, a simple and ancient technology that already shows an advance on chimpanzee skill levels (Whiten et al. 2009a, b). As stone tools became more important in adult lives, children with those adults explored a world increasingly seeded with usable stone, partly made tools and the by-products of tool making. Parental acts bias the environment explored by trial and error learning. As this process elaborated, increasing in fidelity and bandwidth, expertise was acquired and exercised only through very rich environmental support. Those acquiring (say) a craft skill learn by trial and error, but with access to whole, partially complete and failed exemplars of the target artefact, with the aid of tools that are initially chosen by others and with access to raw material in various stages of preparation. They would have many opportunities of observational learning; they would often have access to advice and demonstration, even in those cultures with little explicit teaching. Their trials are guided: Tasks are assigned appropriate to their skill level, but which establish capacities which are platforms for further improvement.

This form of social learning is made possible by a natural bargain. Skilled practitioners ease their own burdens by having apprentices do low to medium skilled work. Apprentices are gofers. They are expected to fetch and partially prepare raw material and to clean and maintain tools and worksites (Stout 2002). This is grunt work from the perspective of the skilled, but for the beginner, it builds basic skills.

Beginners pay for access to expertise and its product by work, while at the same time building their foundational capacities. In short, skill develops in a world not just rich in relevant information but adaptively organised to make that information available and salient. Cognition is environmentally supported. Over both evolutionary and developmental time frames, inner mechanisms have coevolved with and adapted to this rich environment. Language and arithmetical notation enhance our capacity to think, even when we do not have external resources to hand (for example, in doing mental arithmetic).

There is obviously convergence between this evolutionary model of environmentally scaffolded intelligence and the extended mind perspective of Andy Clark and others, that is, with the idea that human cognitive systems contain components that are external to the agent's body and that these are often critical in explaining intelligent action. This convergence is especially marked in more recent defences of the idea. As is well-known, Clark and Chalmers originally developed their case for the extended mind via the example of Otto, an agent with memory failures who compensated for his failing internal memory by keeping information in a notebook he kept to hand. Clark and Chalmers argued that items in the notebook were (amongst) Otto's memories (Clark and Chambers 1998), and their argument depended on a "parity principle". Roughly, if an external resource plays the same functional role in supporting action as an action-supporting internal resource that is uncontroversially cognitive, then the external resource is part of the cognitive system of the agent. Otto's notebook contains information about the location of a museum, playing the same functional role for Otto as ordinary belief plays in the life of agents with normal memory. So it is a belief.

This celebrated example and its associated defence via the parity principle really does the extended mind model no favours. First, it invites a series of sceptical suggestions on the functional difference between Otto's notebook and internally represented information: For example, the notebook is subject to other-party interference and manipulation, or that it is accessed only via other intentional states (beliefs about the content of the book, for example). External representations are stable, physically discrete, assessed through perceptual systems and have representational properties that depend partly on convention. So there are bound to be functional differences between the inner and the outer. Of course, as Clark and Chalmers note, functional similarity does not mean identity in every detail. So they offer a coarse-grained functional analysis that captures (they suggest) both routine cases of dispositional belief and Otto's notebook cases.[3] However, there is a price to going coarse-grained: the less you take internal and external resources to have in common, the less rich and powerful your unified theory of those resources. Much more importantly, though, the criterion of functional similarity threatens to make convincing extended mind cases exceptional, rather than seeing them as central to explaining the problem solving competence of ordinary humans in many circumstances. For example, parity supported cases do not plausibly generalise to other intentional or cognitive states, particularly those with an affective or motivational elements. The more one thinks cognition is embodied; the less one will accept functional equivalence between inner and outer. So, for example, it is hard

[3] Thus, they require that both the resource itself and its informational content be reliably accessed and at low cost and that, as a default, the agent unreflectively trusts the source.

to credibly imagine Otto keeping his preferences in his notebook, representing the information that he is gay, or that he likes blonds. The notebook might of course be an external cue, a prompt that allows Otto better access to his internal, embodied wants and desires. But it cannot substitute for those internal states, for these have a phenomenological, embodied component. In contrast, the notebook can guide Otto to the museum even if he has no moment of recognition when he looks up the address. The notebook might be an external belief store, but not an external store of lusts, longings, hopes and preferences. At most, it gives Otto a new way of auto-prompting his systems of motivation and desire.

John Sutton makes this shift away from the parity principle explicit (Sutton 2010). As he rightly emphasises, the real value of these tightly coupled external resources is that they are functionally distinct from, but complementary to, internal resources (see also Clark 2007, 2008; Menary 2010). Extended minds are more powerful not just because they are bigger, having external as well as internal resources. They are more powerful through differentiation and the division of labour. Thus, Sutton, building on Merlin Donald, points out that externally encoded information is often discrete, coded in representational vehicles which are stable and which are modality and task independent. Memory and other forms of internal representation often do not have these characteristics. Seen this way, we all have extended minds. Cognitive extension is not restricted to external prosthetic substitutes for now degraded internal resources.

So there is in the offing two closely related ideas. The scaffolded mind hypothesis proposes that human cognitive capacities both depend on and have been transformed by environmental resources. Often these resources have been preserved, built or modified precisely because they enhance cognitive capacity. The extended mind hypothesis proposed that human cognitive systems include external components. These components are coupled to human bodies, but not located within human bodies. According to this picture, my glasses are a proper part of my visual system (as a birder, perhaps even my binoculars are, since I carry and use them so regularly). Otto's notebook is a proper part of his memory; an old-style engineer's slide rule is a proper part of the cognitive subsystem that supports his mathematical competence.

Given the obvious affinities of these two ideas, it is natural to wonder whether the scaffolded mind hypothesis and the extended mind hypothesis are causally and explanatorily equivalent. Of course, even if they are causally equivalent, one might offer heuristic advantages. In his *The Extended Phenotype*, Richard Dawkins accepted that individual selectionist models are equivalent to gene selectionist models, but still argued that the gene's eye perspective offered heuristic advantages. In his 2007, Andy Clark runs a similar idea. As he sees it, we can describe environmentally supported thinking using a framework that sees the mind as the biological brain, but which sees that brain as richly interfacing with external factors. But he thinks it is more insightful to see the mind as extended. That extended framework alerts us to the presence and importance of environmental factors and erases the temptation to lapse into Cartesian errors about the nature of the mind.[4] We are inoculated against the temptation to think there is a centre of cognition and consciousness, the location of all real cognition. No one will be tempted to think we

[4] Here Clark has clearly forgotten his Wallace; Chalmers is an unreconstructed Cartesian if ever there was one.

have direct and infallible access to our own states of mind if some of these are written in a notebook. In the next three sections, I pursue this issue of equivalence by identifying and taxonomising different kinds of external resource and hence the different ways external resources contribute to cognitive competence. As a consequence, I shall suggest that the scaffolded mind and the extended mind hypotheses are not equivalent. Rather, the most compelling and plausible cases for the extended mind hypothesis are limiting special cases of scaffolded minds. If that is true, the extended mind picture is not heuristically helpful. It obscures rather than highlights both the continuities and the differences amongst external resources and their contributions to cognitive competence.

In his *Darwinian Populations and Natural Selection*, Peter Godfrey-Smith introduces a very useful analytic tool that I shall co-opt (see Godfrey-Smith 2009). One of his goals is to develop a general characterisation of evolutionary regimes, one that is neutral about the physical details of the evolving agents and their interactions. Instead of asking whether a population of agents evolves under selection, Godfrey-Smith identifies features of agents and populations that come in degrees and which are important to the existence and nature of evolutionary dynamics (for example, the fidelity of inheritance). The result is a way of locating populations in a multi-dimensional space; each variable feature becomes a dimension of that space. Marginal cases of naturally selected populations (say, ones in which success is due to luck rather than fitness) will be close to an edge in this space. This analytic framework has important benefits. First, it does not force categorical distinctions on us if none is to be found in nature. Second, it focuses attention on the right questions: on identifying the features that matter, on the extent to which they vary together and on their dynamics. I propose to use a broadly similar analytic framework in thinking about extended mind cases. Environmental fuels for cognition vary in three dimensions (at least).[5] Each dimension corresponds to a functional relationship between a resource and an agent, and in calling these relationships dimensions, I mean to imply that the relationship is a matter of degree. If this is right, there is a multi-dimensional space of environmental resources which amplify agent cognitive capacity, and the extended mind hypothesis does not make that fact vivid.

Trust

One dimension is trust. Trust involves the agent's assessment of the reliability of their access to a resource and the reliability of the resource itself. Some cognitive resources are automatically trusted; others are used in a more guarded way. The more agents trust a resource, the less they will see themselves as needing redundancy or other insurances against failure. In an earlier response to the extended mind hypothesis, I argued that external resources—like Otto's notebook—are functionally

[5] John Sutton also develops a dimensional analysis of environmental resources, but his dimensions are differences in material substrate rather than in functional relations between agent and resource. Thus, he distinguishes between artefacts, the use of other agents and the reliance on stable features of the physical environment (Sutton 2006).

different from internally stored information. They are not trusted, or should not be trusted, to the same extent that we trust internal resources (Sterelny 2004). For storage in a world of only partially cooperative agents made Otto's exomemories vulnerable to confiscation, corruption and deception. Otto's competitors have the opportunity to steal his notebook, erase passages in it and add deceptively to it. If Otto is rational, he will be aware of such a danger and will be wary of committing himself to a high-stakes action on the basis of his notebooked beliefs alone. Likewise, Leonard, from the film Memento, had best avoid binge drinking or sleeping pills, else his hunt might be subverted by others adding misleading tattoos to him in his sleep. Our ordinary, internally stored memories are not vulnerable in the same way.

Once the case for an extended mind hypothesis comes to depend on the complementarity of internal and external resources in an integrated system, the functional similarity of inner and outer resources is not the key issue. Moreover, Clark, in response, has rightly pointed out that the initial contrast was overdrawn. It exaggerates the security of internal resources. And though we typically have the capacity to sceptically assess external informational resources in our environment, de facto we often treat such resources as trustworthy. After all, attention and scrutiny are scarce resources, to be used only when necessary (Clark 2010).

Moreover, my earlier ideas on the importance of contested space were over-influenced by Machiavellian models of social interaction. Hostile manipulation of my informational environment is a serious danger, but only in a restricted range of interactions, for example, in one-on-one high-stakes negotiations. To recycle an example from another context, Monty Python's celebrated sketch of a deceptive Hungarian–English phrasebook and the catastrophe that befalls its innocent Hungarian user is enjoyable in part because it is difficult to envisage circumstances in which an author would gain from producing a maliciously misleading phrasebook, for an author cannot know when, where, by whom or with what effect such a book will be read. Thus, informational resources in a shared space are sometimes reliable *because* they are shared. Such resources are the joint product of many agents and are typically used at unpredictable times and places. In many circumstances, public domain resources cannot safely be used to manipulate a specific target for a specific purpose. While it might be possible to derail Otto's social plans by erasing a crucial line in his notebook, it would be much harder to do so by, say, corrupting the tools he uses to negotiate the subway system. Maps of the underground exist in many copies, with many users able to spot fakes. It is a dispersed, highly redundant resource. Even if, improbably, a manipulative enemy of Otto succeeded in replacing all the subway maps at Arsenal tube station (his normal entry point, let's suppose, to the system), every carriage of every train has another. We do not have to worry about deceptive underground maps: We can trust both the resource and our access to it. Single-sender, single-receiver systems are quite different: It would be unwise to treat the vendors' private description of the condition and value of their house with unreflective trust.

There are of course examples of dispersed, many-copy, multi-user informational vehicles that are deceptive and manipulative. Propaganda and advertising provide examples. But such cases of outright deception are rare. Reliability is a more pervasive issue. Agents can have more or less reliable access to a resource, and the

resource itself can be more or less reliable. But agents can also manage reliability, often learning rules of thumb about access and reliability. Canberra, for example, is liberally provided with external navigational resources in the form of road signs. These cue names, directions to important locations, speed limits and rights of way. They are not deceptive; they are regularly present; their content is highly reliable. This set of resources scaffolds navigation around Canberra's confusing street network. My birding e-mail list, birding.oz, on the other hand, is significantly less trustworthy. I typically do not have access to it when I need it most, which is when I am in the field. And while actual deception is extremely rare, I suspect self-deception and mere error are not. There are a few experts of very high repute that post to it. But unless one of these vindicates a report of a rarity, I would be cautious before galloping off in search of a special bird on the basis of a report on the list. I use it in planning trips, but with caution: There is a good deal of noise in the signal. It is a valued resource, but I do not use it in the way I use my well-practiced cognitive map of Wellington, or as Otto uses his notebook. While the well-adapted agent will certainly use less trustworthy resources, it is risky to become systematically dependent on such resources (of course that might be the least bad option).

Interchangeability, individualisation, entrenchment

Let us return to the extended stomach. I have a decent set of knives and a superior set of choppers. But they are standard, shop-bought instruments. Others have identical implements, and from time to time, other cooks have used my knives and choppers. These are my implements for my use—they are not shared resources—but because they are standardised, they are interchangeable. I can use others' knives without disrupting my cooking routines, and likewise my friends can use mine. Contrast a professional cook, who will certainly not let anyone else near his or her knives. Professional cooks' knives are often significantly individualised. They are not just chosen with a specific eye to individual skills, preferences and needs. They can be retouched for their user, sometimes significantly. They will certainly be honed to individual taste. Sometimes the handle and even the blade are modified, so that weight, balance, length and grip fit the cook.

Modification will often be mutual. The cook's physical routines and skills adjust to the weight, balance and sharpness of her knives; the cook's movements are adapted to her knife, just as the knife has been individualised to her prior preferences. If forced to use different knives, an expert will cope, but physical routines will be less comfortable and precise. The same is true of many highly skilled activities using individualisable artefacts. An expert batsman, for example, will choose a bat with a weight, balance and flex that suites him, but he may well modify the grip (adding, for example, an extra layer of rubber). But as he uses the bat, over time his own physical skills will in turn be tuned to its profile; if forced to play with another, it may well seem clumsy, rigid or insubstantial. The tool is not just individualised, it is entrenched. It has been claimed that to those long blind, their stick feels like an extension of their hand, and for them, phenomenologically the interface between body and world seems to be at the end of their stick rather than at

the end of their hand.[6] If so, I conjecture that this is because the stick is both individualised and entrenched. Substitute a stick of different weight, balance, length or flexibility, and the interface will shift inward.

This gradient from standardised and interchangeable resources to individualised and entrenched ones applies to cognitive resources too. Most of the books in my professional library are interchangeable. They can and do offer my graduate students the same support that they offer me. But there is a small subset on which I have seriously worked, and these are massively individualised with underlining, marginal comments and post-it notes with longer commentaries. The text of my first copy of *The Extended Phenotype* has disappeared under a mould of marginalia, of hooks to my own thought and work. I doubt that any of these professional resources have become entrenched; I do not think that my professional routines and working habits have become adapted to specific texts, as I adapt them to me. No single work is sufficiently salient. But I suspect that some historians of philosophy, who live their professional lives with just a few texts, adapt to their canonical texts. To a Locke scholar, the cadence, rhythm, balance and vocabulary of seventeenth century prose can come to seem natural and transparent; that of the early twenty-first century tangled and jargon-ridden. Their skill set and working routines adapt them to a few particular resources. Individualising resources and entrenching them takes time: The relationship between agent and environmental support develops over time. Once developed, it is stabilised. So such individualised, stabilised resources will be persistently used. Clark resists using persistence as a criterion which distinguishes external cognitive components from mere environmental resources (Clark 2007, 2008). But Otto's notebook and similar examples are used persistently because they have been individualised. They have value added by being both adapted by the agent and by the agent adapting to them.

For the craftsman and the cricketer, their special tools feel like an extension of their body. Likewise, the extended mind framework seems most natural with highly individualised and entrenched cognitive resources. Consider, for example, the support offered by other agents. Other agents are often importance resources for our cognitive projects: cueing, demonstrating and advising. But even those sympathetic to the extended mind model rarely treat other agents as part of an extended mind (though there are partial exceptions; Clark and Chambers 1998). Yet to young children, mothers are like Otto's notebook to Otto: reliably and easily available, routinely used, substituting for imperfect memory and trusted by default. Likewise, the division of cognitive labour is of central importance in explaining both the acquisition and the exercise of many cognitive competences. Many academic projects depend on collaboration and on technical and specialist support. When confronted with new problems—for example, an intractable program on an upgraded computer—I will, as Clark says, "soft-assemble", on the fly, an appropriate though often ephemeral set of resources. These often include other agents. But while crucial, these agents are not my resources alone, and they resist individualisation. I sometimes have access to what others know and sometimes even trust that access and the information to which I have through that access. But I cannot adapt their minds to my purposes, not in a permanent, sustained and reliable way. I cannot scribble marginalia in my collaborator's belief box,

[6] I was alerted to this idea in Malafouris (2008), but I have learned from Richard Menary that it dates back to Mearleau-Ponty and that it has been defended in some detail in Gallagher (2005).

highlight the important points or delete material no longer relevant. Five-year-olds do not individualise their mother, as Otto does his notebook.

The individual and the collective

In his paper in Richard Menary's collection, John Sutton discusses a terrific example of environmentally supported cognitive capacity (Sutton 2010). It turns out that historians of English drama have been puzzled and impressed by the repertoire size of Elizabethan theatre companies. They had small troupes of players and large, rapidly changing and rarely repeated repertoires. A working actor might have to master 70 roles a year, perhaps 50 of which would be new, this while acting in six different plays a week (Tribble 2005). That raises a puzzle. How did these small troupes memorise so many and such rapidly changing roles? As Sutton tells it (drawing on the work of Evelyn Tribble), both the scripts and the physical layout of the theatre itself were organised to constrain the range of possibilities and to cue action at the right moments. Plays themselves had a somewhat stereotyped overall organization, and the stage had a stable organization through which action was funnelled in highly constrained ways. Likewise, scripts provided actors with skeletal information on how a given play varied from the predictable flow of entry and exit. There was no magic bullet that explained actor capacity to master many roles rapidly. Rather, many aspects of the physical organization of the theatre itself, the thematic organization of the play, the minimal stage directions cueing a few but significant moments and the responses of the rest of the troop all work together to support the actor's mastery of their roles.

While this is a convincing example of the environmental support of cognitive capacity, it is hard to shoehorn it into the extended mind model. As Sutton himself notes, what is supported is a collective activity. Moreover, that collective activity is supported by collectively made and used props and other resources. An Elizabethan theatre troop is a collective: a set of agents acting jointly and collaboratively depending on one another. They take advantage of a space they have in part inherited from others. But they jointly organise that space and use stage directions that have been adapted to all of them, rather than to any single one of them. Menary (2007, 2010) and Sutton (2010) distinguish cognitive components from mere resources by appeal to the idea of an integrated system. But while there may be an integrated system here, it does not consist of an agent plus a set of external components, chosen for and adapted to their role in supplying that agent with a grip on his part(s). These are multi-part, multi-actor plays, and this transient assembly of resources needs to support them all. Each is both supported by this system and contributes to it, helping others situate themselves.[7]

[7] Of course it is possible to suggest that there is a collective agent, the cast itself, that knows the script and in virtue of this array of environmental resources. But while there may be a collective agent, there are also individual agents, each of whom knows his part, though only in context and with the aid of various forms of support. Little of that can be regarded as part of an agent's cognitive system, inter alia, because agents would have to be part of one another. Mark Rowlands, in distinguishing external components of the mind from mere resources, appeals to an intuitive notion of "belonging to" an agent (Rowlands 2009). None of the crucial resources here intuitively belongs to any specific agent.

In contrast, Otto's notebook, my reading notes on Dawkins or a birder's field sketches and notes at an unidentified sighting are individually produced and used. The poster examples in the extended mind literature have typically been individually used props: paper drafts, a mathematician's working notes on paper or whiteboard and an architect's first sketches. These examples can be misleading, as human problem solving activity is often social and much more dependent on communal resources. In this respect, a team of scientists working together in a lab is a more typical example of human problem solving than my revising drafts alone in my office (Sterelny 2004).

The distinction between individual and collective resources is important, for collective resources have distinct individual and intergeneration dynamics. Defenders of the extended mind hypothesis sometimes write of language as the "ultimate artefact" (Wheeler 2004 building on Clark 1997). It is right to identify language as one of the great transformers of the human lineage. It is an immensely important resource for formulating and communicating thought; almost certainly, our language-using habits have transformed the internal processes of human minds both on evolutionary and developmental time scales (Deacon 1997). We adapt the expressive powers of language to our own purposes, but no doubt we have also adapted to it. But in its transformative power, language is by no means unique. In other work, I have argued that the socially organised transmission of expertise has deep roots and has equally transformed human lives and minds. The importance of expertise has bought with it selection for social learning and teaching (demonstration, imitation, theory of mind), selection for a more broadly cooperative social life. This form of social learning has contributed to the evolution of the distinctively human life history, with extended childhoods, long lives and intergenerational resource transfers (Sterelny 2007, 2010). More recently, the development of mathematical notation and depictive representation—maps and diagrams—has enhanced our capacities for quantitative reasoning and complex physical construction. Writing, likewise, reduces burdens on memory dramatically and further decouples language from immediate context. The plasticity of the human brain implies that the routine use of these tools has internal consequences, consequences for our inner representational capacities. These are most unlikely to be environmental supports whose use leaves the inner machine as before.

Unlike an architect's preliminary sketches, my annotated Dawkins or Otto's notebook, these transforming cognitive technologies resemble termite mounds, beaver dams, gerygone nests and other classic extended phenotype traits. Almost certainly, these technologies evolved by cumulative trial and error, as do those complex extended-morphological adaptations. But the mechanism of inheritance is cultural rather than genetic, and this difference leads to another. Inheritance is not strictly vertical, from parent to offspring. Rather, it is often oblique: Information flows from many members of the parental generation. We do not learn language *only* from our parents, and very likely the same is true of many forms of expertise and of the skills of literacy, numeracy and depictive media.

These environmental resources structure and amplify cognition, as children develop into mature members of their community. But at least initially, those supports are provided and shaped by others, by the previous generation. The previous generation acting collectively are children's most important environmental

resource. That generation shapes the next both directly, through implicit and overt teaching, and indirectly, through their modification of the environment children explore in trial and error activity. I alone shaped my copy of *The Extended Phenotype*, as a professional resource for my research. But my language, my tools of quantitative reasoning, the norms of social interaction through which I read others and are read by them were shaped by others for me.

In his 2007, Andy Clark noted that an organic brain chooses and assembles the resources that make much problem solving possible. True. But in many critical cases, those brains belong to members of the previous generation, not to the agent faced with the problem. Not many of us design our own mathematical notation. Thus, the cognitive competence of generation $N+1$ individually and collectively depends on cognitive provisioning by generation N. The most critical, mind-and-brain-shaping environmental supports for cognition are these cumulatively built, collectively provided tools for thinking, tools that are provided to many or all of a generation by many or all of the previous generation. They do not naturally fit an extended mind framework. The same is true of the extended stomach: The pervasive changes in life history and morphology result not from my individualised and privately owned cooking equipment but from cumulatively built and collectively transmitted changes in the way we produce and prepare food. Cooking, too, is an "ultimate artefact".

Finale

Andy Clark accepts that it is possible to treat external supports for cognition as scaffolding, as showing that humans have their cognitive capacities in part through the way they are embedded in their environment. But he thinks it is more illuminating and insightful to see some resources as genuine constituents of human minds (Clark 2007, 2008). He argues that the extended mind framework is more heuristically compelling. He takes this view in part, I think, because he thinks alternative frameworks simply understate the importance and the transformative power of external components.

However, his line of argument also depends on treating the extended mind and embodied mind frameworks as if they are a package deal. For example, one key example is a detailed defence of the importance of gesture. He shows that gesture is an aid to thought, not just communication, because gesture is itself a form of mental representation that supplements and complements other media. It is hard to evaluate the overall cases for an embodied mind framework, in part because the commitments of that framework are rather elusive. One element seems to be the idea that in perception, action and motivation, structures in the body support intelligent action without themselves being representations. There are elegant examples, especially from acoustic perception in animals, that show how structures in the ear support acoustic discrimination of location by modulating the time of arrival of sound to the two ears, structures which do not work by representing location. Clark's mole cricket example is particularly impressive. Likewise, it is likely that arousal and mood profoundly affect motivation and hence action, without representing anything in body or world. The extent of such effects on human action remains unclear. But it is

plausible that there are structures in the body that substitute for representation and computation. In any case, the nervous system is distributed throughout the body, and so I agree that our cognitive engine is the brain situated within our body and using its resources, not just the brain itself. However, one can buy the extended mind hypothesis without the embodied cognition framework and vice versa. These frameworks are not a package deal; each can be accepted without the other. So even a compelling case for embodied cognition does not give us a reason to prefer the extended mind framework to the alternative scaffolding model.

So I think the heuristic considerations point the other way. The framework outlined in this paper is an existence proof that a theory of scaffolded cognition need not understate the importance of the environment. It is true that the extended mind model seems very plausible when we focus on highly trusted, individualised and entrenched, single-user resources. But that is one corner in a 3D space of environmental scaffolds of cognitive competence and that in turn is just one case of the environmental support of adaptive phenotypes. No clear error is made in reserving a special label for this region of space, though obviously the boundary between external components of the agent's mind and mere resources for that mind must be arbitrary. But I do not see how privileging that region would be helpful. It obscures the fact that extended mind cases are special cases of a general phenomenon. If I am right about language and the like, the most crucial environmental supports for cognition do not lie within the extended mind region. Moreover, reserving a special label for the extended mind cases does not focus on the right issues, for the taxonomy of resources I have developed raises further questions. Are these the only dimensions of importance? To what extent are they independent? What are the dynamics of movement in the space? Under what circumstances, for example, do collectively used resources become segmented into single-user resources and vice versa? How do resources become individualised and entrenched? How do actual levels of trust covary with trustworthiness? It is much easier to see these questions with a model that focuses on the space itself. I see no compensating advantage with the alternative framework.[8]

References

Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge: MIT.

Clark, A. (2007). Curing cognitive hiccups: A defense of the extended mind. *Journal of Philosophy, 104* (4), 163–192.

Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford: Oxford University Press.

Clark, A. (2010). Memento's revenge: The exended mind extended. In R. Menary (Ed.), *The extended mind*. Cambridge: MIT.

Clark, A., & Chambers, D. (1998). The extended mind. *Analysis, 58*(1), 7–19.

Dawkins, R. (1982). *The extended phenotype*. Oxford: Oxford University Press.

Dawkins, R. (2004). Extended phenotype—but not too extended. A reply to Laland, Turner and Jablonka. *Biology and Philosophy, 19*, 377–396.

Deacon, T. (1997). *The symbolic species: The co-evolution of language and the brain*. New York: Norton.

[8] This paper benefited greatly from generous feedback from three audiences. These were the 2009 New Zealand Division of the Australasian Association of Philosophy Conference in Palmerston North; Embodied Cognition, Enactivism and the Extended Mind (a Wollongong workshop) and the ANU philsoc.

Dennett, D. C. (2000). Making tools for thinking. In D. Sperber (Ed.), *Metarepresentation: A multidisciplinary perspective* (pp. 17–29). Oxford: Oxford University Press.
Foley, R., & Gamble, C. (2009). The ecology of social transitions in human evolution. *Philosophical Transactions of the Royal Society B, 364*, 3267–3279.
Gallagher, S. (2005). *How the Body Shapes the Mind*. Oxford: Oxford University Press.
Godfrey-Smith, P. (2009). *Darwinian populations and natural selection*. Oxford: Oxford University Press.
Gurven, M., & Hill, K. (2009). Why do men hunt? A reevaluation of "man the hunter" and the sexual division of labor. *Current Anthropology, 50*(1), 51–74.
Hill, K., & Kaplan, H. (1999). Life history traits in humans: Theory and empirical studies. *Annual Review of Anthropology, 28*, 397–430.
Laland, K. (2004). Extending the extended phenotype. *Biology and Philosophy, 19*(3), 313–325.
Laland, K. (2007). Niche construction, human behavioural ecology and evolutionary psychology. In R. Dunbar & L. Barrett (Eds.), *Oxford handbook of evolutionary psychology* (pp. 35–48). Oxford: Oxford University Press.
Laland, K. N., Odling-Smee, J., & Feldman, M. (2000). Niche construction, biological evolution and cultural change. *The Behavioral and Brain Sciences, 23*, 131–175.
Malafouris, L. (2008). Beads for a plastic mind: The 'blind man's stick' (BMS) hypothesis and the active nature of material culture. *Cambridge Archaeological Journal, 18*(3), 401–414.
Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Basingstoke: Palgrave Macmillan.
Menary, R. (2010). The extended mind and cognitive integration. In R. Menary (Ed.), *The extended mind*. Cambridge: MIT.
O'Connell, J. F., Hawkes, K., & Jones, N. B. (1999). Grandmothering and the evolution of Homo erectus. *Journal of Human Evolution, 36*(5), 461–485.
Odling-Smee, J., Laland, K., & Feldman, M. (2003). *Niche construction: The neglected process in evolution., monographs in population biology*. Princeton: Princeton University Press.
Rowlands, M. (2009). Extended cognition and the mark of the mental. *Philosophical Psychology, 22*(1), 1–19.
Serventy, V. (Ed.). (1982). *The wrens and warblers of Australia, National Photographic Index of Australian Wildlife*. Sydney: Angus and Robinson.
Sterelny, K. (2003). *Thought in a hostile world*. New York: Blackwell.
Sterelny, K. (2004). Externalism, epistemic artefacts and the extended mind. In R. Schantz (Ed.), *The externalist challenge. New studies on cognition and intentionality* (pp. 239–254). Berlin: de Gruyter.
Sterelny, K. (2006). Cognitive load and human decision, or, three ways of rolling the rock up hill. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Volume 2: culture and cognition* (pp. 218–233). Oxford: Oxford University Press.
Sterelny, K. (2007). Social intelligence, human intelligence and niche construction. *Proceedings of the Royal Society London (series B), 362*(1480), 719–730.
Sterelny, K. (2009). Novelty, plasticity and niche construction: The influence of phenotypic variation on evolution. In A. Barberousse, M. Morange, & T. Pradeu (Eds.), *Mapping the future of biology. Evolving concepts and theories*. Dordrecht: Springer.
Sterelny, K. (2010). From hominins to humans: How sapiens became behaviourally modern. *Philosophical Transactions of the Royal Society, London, B*, in press.
Stout, D. (2002). Skill and cognition in stone tool production: An ethnographic case study from Irian Jaya. *Current Anthropology, 43*(5), 693–722.
Sutton, J. (2006). Distributed cognition: Domains and dimensions. *Pragmatics & Cognition, 14*(2), 235–247.
Sutton, J. (2010). Exograms, interdisciplinarity and the cognitive life of things. In R. Menary (Ed.), *The extended mind*. Cambridge: MIT.
Tribble, E. (2005). Distributing cognition in the globe. *Shakespeare Quarterly, 56*(2), 135–155.
Wheeler, M. (2004). Is language the ultimate artefact? *Language Science, 26*(6), 693–715.
Whiten, A., McGuigan, N., Marshall-Pescini, S., & Hopper, L. M. (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society, series B, 364*(1528), 2417–2428.
Whiten, A., Schick, K., & Toth, N. (2009). The evolution and cultural transmission of percussive technology: Integrating evidence from palaeoanthropology and primatology. *Journal of Human Evolution, 30*, 1–16.
Wilson, R. A. (1994). Wide computationalism. *Mind, 103*, 351–372.
Wrangham, R. W. (2001). Out of the pan, into the fire: How our ancestors' evolution depended on what they ate. In F. de Waal (Ed.), *Tree of life* (pp. 121–143). Cambridge: Harvard University Press.
Wrangham, R. (2009). *Catching fire: How cooking made us human*. London: Profile.
Wrangham, R., Jones, J., Laden, G., Pilbeam, D., & Conklin-Brittain, N. (1999). The raw and the stolen: Cooking and the ecology of human origins. *Current Anthropology, 40*(5), 567–594.