James Bond and the Barking Dog: Evolution and Extended Cognition*

Lawrence Shapiro†‡

Prominent defenders of the extended cognition thesis have looked to evolutionary theory for support. Roughly, the idea is that natural selection leads one to expect that cognitive strategies should exploit the environment, and exploitation of the right sort results in a cognitive system that extends beyond the head of the organism. I argue that proper appreciation of evolutionary theory should create no such expectation. This leaves open whether cognitive systems might in fact bear a relationship to the environment that leads to their extension.

1. Introduction. Appeals to evolutionary theory have played a supporting role for the idea that cognition extends beyond the boundary of the head. At its sketchiest, the reasoning is this: natural selection is likely to favor cognitive strategies that exploit environmental structure in some special way; when these strategies exploit this structure in the right way, cognition extends into the environment. As I said, this is sketchy. Left to be filled in are details about what the "special" and "right way" of exploitation is and what it means to say that cognition extends. The dispute between the most prominent champions of extended cognition (Andy Clark, Mark Rowlands, and Rob Wilson) and its fiercest critics (Fred Adams and Ken Aizawa) is largely over these details. For instance, Clark is quite clear that in those cases when cognition does extend into the environment there exists a particular sort of coupling between the brain and the environment, one in which the environment enhances or serves an information-pro-

*Received July 2009; revised January 2010.

†To contact the author, please write to: Department of Philosophy, 5185 Helen C. White Hall, 600 North Park Street, University of Wisconsin, Madison, WI 53706; e-mail: lshapiro@wisc.edu.

‡Elliott Sober provided me with helpful comments on an earlier draft of this article. I am grateful also to Fred Adams, Ken Aizawa, John Bickle, Carrie Figdor, Carl Gillett, Tom Polger, and Jackie Sullivan. Comments from an anonymous referee for this journal were especially valuable.

Philosophy of Science, 77 (July 2010) pp. 400–418. 0031-8248/2010/7703-0003$10.00
Copyright 2010 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

cessing role (2008).[1] Similarly, for Rowlands, cognition extends when it involves the manipulation of *information-bearing pieces* of the environment. In turn, Adams and Aizawa attack proponents of extended cognition for, in their view, failing to identify a relationship of environmental exploitation that does more than simply make the environment a causal contributor to, rather than constituent of, cognitive processes. Yet, lost in the dust of this conflict is a careful examination of the first step in the reasoning I sketched above. Should we expect organisms to have evolved cognitive strategies that exploit the environment?

Clark's *007 principle* and Rowlands's *barking dog principle* answer this question affirmatively. The 007 principle especially has been greeted with general approbation among defenders of extended cognition (Wilson 2004; Menary 2007) and receives frequent reaffirmation from its author (Clark 1997, 2008; Wilson and Clark 2009). In the first part of this article, I shall explain Clark and Rowlands's principles and briefly examine how they might serve to justify the thesis of extended cognition. I note at the outset that there are various arguments from natural selection's purported preference for cognitive strategies that exploit the environment to the conclusion that cognition extends. Some (e.g., Rowlands) seem to see the connection as being more direct than others (e.g., Clark). However, the main business of this article unfolds in the middle sections, where I argue that the 007 principle and the barking dog principle are less obviously true than their authors believe. This sets the stage for the concluding section, where I consider more broadly the strategy of deriving a theory of mind from evolutionary considerations.

2. From Environmental Exploitation to Extended Cognition. Clark's 007 principle asserts that "in general, evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand-in for the information-processing operations concerned. That is, know only as much as you need to know to get the job done" (1989, 64; see also Clark 1997, 46; Wilson and Clark 2009).[2] In similar spirit, Rowlands proposes the barking dog principle: "If it is necessary for an organism to be able to perform a given adaptive task T, then it is differentially selectively *dis*-advantageous for that organism to develop internal mechanisms sufficient for the performance of T when it is possible for the organism to perform

1. For discussion of what this right sort of exploitation involves, see (Clark and Chalmers 1998; Clark 2008, 41, 96, 126, 131).

2. The principle receives its name because, as a member of Her Majesty's Secret Service, James Bond is told only what he needs to know in order to complete his mission.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

T by way of a combination of internal mechanisms and manipulation of the external environment" (1999, 80).[3]

Because Clark's principle concerns information-processing operations in particular and Rowlands's makes a claim about adaptive tasks in general, the 007 principle is a special case of the barking dog principle. On the assumption that cognition involves information processing, both principles imply that natural selection should have favored cognitive systems that exploit, manipulate, or otherwise take advantage of the environment in ways that minimize the load on internal processes.

The guiding intuition that takes one from evolution to extended cognition is this. Imagine that a given task requires a fixed quantity of resources for its completion. Suppose that there are two strategies that might work to satisfy the task demands. Strategy 1 *mines* the organism, meaning that most of the resources required for the strategy come from the organism itself. Strategy 2 mines the environment: it draws its resources primarily or substantially from outside the organism. But, the thinking goes, organism-mining strategies come at greater cost to the organism than environment-mining strategies, and hence organisms that adopt the first kind of strategy are at a disadvantage when competing against those deploying the second.

This intuition, in the context of psychology, seems to lend support to the idea that natural selection will favor cognitive capacities that "offload" their burdens onto the environment. As a simple example, consider a strategy for calculating the product of three-digit numbers. One that mines only internal cognitive resources will be more costly to an organism than one that mines the environment by, for instance, off-loading steps in the multiplication task to pencil and paper or a calculator. Similarly, memorizing a long list of phone numbers carries greater internal costs than writing these numbers down in a business diary.

Clearly, these points on their own do not demonstrate that cognition actually extends into the environment. As I mentioned above, this conclusion requires further discussion of (i) the nature of the relationship between the internal processes and the external structures they exploit and (ii) the claim that cognition itself extends into the environment. Regarding the first point, Adams and Aizawa (2001, 2008), Rupert (2004), and Block (2005) have charged advocates of extended cognition with confusing causal contributors to a process with constituents of a process. Granting that cognitive systems exploit the environment, these critics then ask why this should entail that the environment becomes part of the system

3. The name of this principle derives from the reasonable suggestion that if you have a dog, then you do not have to bark yourself. The idea is that if there is something in the environment that can do the job for you, then you do not have to do it yourself.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

rather than simply contributing to the system's success. As Shapiro and Spaulding (2009) summarize the criticism, "use of a hammer, after all, does not entail that the hammer becomes literally part of the arm."

Menary (2007), despite his own commitment to extended cognition, appears to think that this criticism can be made to stick against Clark and Chalmers (1998).[4] Citing Clark and Chalmers's example of the scrabble player who rearranges scrabble tiles in the course of "thinking up" new words, Menary describes their argument for extended cognition this way: "C&C take the rearrangement of scrabble tiles on a tray . . . *to be part of thought*. . . . This is because the scrabble tiles are external vehicles and their manipulation is part of an extended cognitive process" (2007, 49). However, Menary then notes, critics like Adams and Aizawa can easily dismiss Clark and Chalmers's conclusion because they can "accept that there are external resources to aid cognition, as well as a class of actions which have the purpose of simplifying a problem—such as the reorganising of the scrabble tiles. This leaves them free to deny that there are any good grounds for identifying them as cognitive" (49).

Yet, almost certainly Menary has misinterpreted Clark and Chalmers's argument. First, as he states it, the argument clearly begs the question. Clark and Chalmers must convince us that use of the scrabble tiles is part of an extended cognitive process, and thus they cannot, as Menary says, base their case for cognitive extension on the grounds that the manipulation of scrabble tiles is part of an extended cognitive process. The challenge, of course, is to defend a reason that environmental structures should count as constituents of cognitive processes rather than causal contributors. Toward this end, both Clark and Wilson (Wilson 2004; Clark 2008; Wilson and Clark 2009) have emphasized the significance of a deep functional integration between an agent and external resources. Although the precise nature of this integration is not always clear, an important feature of it seems to be its incorporation of a *loop* (Shapiro 2010). As with a turbo-driven engine that uses the very exhaust it produces to increase its own power, a cognitive system will on occasion structure the environment in a manner that enhances its abilities, which in turn enables it to structure the environment more effectively, enhancing its abilities further, and so on (Wilson 2004, 195; Clark 2008, 131; Wilson and Clark 2009, 67). When cognitive abilities depend on "loopy" interactions with the environment of this sort, the suggestion that external structure is a part of, rather than a mere contributor to, cognition becomes much more palatable.

But the other question to consider concerns the claim that cognition

4. Menary's own case for extended cognition rests on the idea that an agent might be integrated with environmental features in a manner that respects and maintains certain norms. Discussion of this interesting idea would take us too far afield.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

extends. Adams and Aizawa have interpreted this to mean that cognitive processes themselves take place beyond the boundaries of the brain. This invites their jest that Clark would regard a pencil, when coupled to a mathematician, as itself capable of thinking that $2 + 2 = 4$ (Adams and Aizawa 2010). But the claim that cognition extends is ambiguous. It might mean, as Adams and Aizawa suppose, that cognitive processes themselves occur in the environment outside the agent, as it would if the pencil were indeed capable of thought. Alternatively, the claim might mean that cognitive processes are realized to some extent in parts of the world external to the brain, while the brain remains a necessary component in cognitive processing. Or, finally, it might mean that cognitive processes on occasion incorporate parts of the external world, without additionally claiming that cognitive processes take place within these parts. Thus, one might suppose, for instance, that the use of hand gestures is deeply functionally integrated in certain spatial-reasoning tasks, and thus gestures are a part of the spatial-reasoning system, without having to suppose that spatial reasoning itself takes place in the gesturing hands.[5]

I hope to have said enough to indicate that the path from principles like 007 and the barking dog to extended cognition is by no means direct. These principles, if true, suggest at most that natural selection will favor embedded or situated cognitive systems—cognitive systems that entangle themselves in environmental structure. Exactly how embedded cognitive systems develop into extended cognitive systems remains, as we have seen, controversial. Nevertheless, proponents of extended cognition certainly do see this entanglement as a first step on the path toward extension. Without environmental mining of the sort I described earlier, cognition would not extend. I propose for now to leave aside issues concerning the transition from embedded cognitive systems to extended cognitive systems. This is not my primary interest. Rather, I wish to examine the principles that many who favor extended cognition take to justify confidence that cognitive systems will be embedded—the principles that mark the first step on the path toward cognitive extension.

Fortunately, both the 007 principle and the barking dog principle can be evaluated independently of the case for extended cognition. At issue is an answer to this question: Should we expect to find, within a population of organisms that has evolved by natural selection, a tendency toward environment-mining strategies in favor of internal-mining strategies? Of course, insofar as proponents of extended cognition seek support from evolutionary theory, they are unlikely to be pleased with a negative answer.

5. I gather that this is something like Goldin-Meadow's (2003) view.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

3. In Support of the 007 Principle. Clark never provides an argument for his 007 principle. This is not to say that he believes it to be an obvious truth. Rather, Clark finds support for the principle in the pioneering work of roboticists such as Rodney Brooks (1991), whose "creatures" abandon elaborate models of their environments in favor of instantaneous and frequent updates about the surrounding world. In starkest terms, and not with complete fidelity to Brooks's research, the contrast is this. Imagine two robots with different strategies for navigating through a busy environment. The first, Cogito, uses input systems to create a representation of its world. This representation is stored in a location where it is accessible to the several or many modules that control Cogito's behavior. Thus, Cogito's goal module will examine the representation in order to plan a heading that will take Cogito to a desired location. Cogito's object-avoidance module will check the representation to make sure that the heading does not cause Cogito to collide with objects. Other modules consult the representation to ensure that should Cogito stray from its heading, it will find its way back on track and move toward its goal.

Cogito's navigational strategy is cumbersome. First, Cogito must be equipped with a perceptual system adequate to the task. Studies in computer vision have revealed some of the complexities such a system involves. Algorithms for computing shapes and surfaces in the environment must recover and amplify information that comes from shading, texture gradients, and illumination shifts. These algorithms are not simple. Vision researchers remain actively engaged in trying to understand how the human visual system extracts from information on the 2-D retinal surface a perception of a 3-D world. Second, having produced a map of its surroundings, Cogito must have the resources to update it as the world changes around it. Some of these changes will arise through no fault of its own. Objects in an environment may move about, creating obstacles where before there were none or removing obstacles around which Cogito had already planned a course. Other changes arise as Cogito advances toward its goal. Perhaps Cogito's actions cause changes in the environment—as Cogito heads down a hallway, its motion upsets a broom leaning against a wall, which falls in its path. Cogito must be prepared to update its initial model of the world whenever the world changes, but then Cogito must be aware of the changes as they occur, which would seem to require comparing a present representation of the environment to past representations. Finally, Cogito must exhibit versatility. The plan Cogito's goal module originally designed must be sufficiently flexible to accommodate whatever changes take place in the environment. As the environment changes, so must Cogito's plans.

Cogito's competitor, Reflex, goes about its business in a strikingly different manner. Reflex's goal module, like Cogito's, fixes on a heading.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

However, rather than creating a representation of the world that will be of service to a variety of modules that must act in concert to get Reflex from here to there, Reflex has a collection of very simple machines that connect sensing directly to action. Reflex's goal module starts it on a course toward the desired location. A guidance module, equipped with a compass, is on standby to return Reflex to its heading should it stray off course. As Reflex approaches an obstacle, sensors on Reflex's ventral side shut down Reflex's motors. Reflex stops. Reflex does not bother to create a representation of the object in front of it. Its shape and size are unimportant. If it is large enough to trigger Reflex's sensors, that is all that matters. Once stopped, Reflex's avoidance module turns on. The module causes Reflex to turn right. Reflex will now move forward until sensors on its left side no longer detect a nearby surface. When signals from the sensors end, Reflex stops, turns left again, and moves forward until once more Reflex's port-side signals cease. Now Reflex has just one last left turn to make, at which point it moves forward the same number of steps that it has moved to the right. Reflex has moved around the object, and a right turn will put it back on track.

Rather than relying on models of its world to navigate successfully through its environment, Reflex uses the world itself. Whereas Cogito must first create a representation of an obstructing object in order to plot a course around it, Reflex need not bother. Reflex's "awareness" of the world consists in a stream of sonarlike readings bouncing directly from the world. Because Reflex's sensor signals cause action directly, the creation and storage of a representation is unnecessary.[6]

I share Clark's judgment that Cogito's cognitive "style" is more complex or costly or unwieldy than Reflex's. Indeed, let us grant that there is a metric for quantifying the cost of cognitive architectures and that on this metric Cogito's costs more than Reflex's. But from this stipulation, what can we say about the probability that the cognitive style in a population of evolved organisms will more likely be similar to Cogito's or Reflex's? No doubt (we assume) robots like Reflex are more likely to crawl from roboticists' labs than Cogito. But how reasonable is the assumption that evolution will show the same preference? Significantly, as we will see, evolution faces historical constraints that engineers do not. Engineers have the advantage of being able to start their designs from scratch. That evolution cannot do this marks a difference, the importance of which becomes very apparent when analyzing Rowlands's argument for his version of the 007 principle, that is, the barking dog principle.

6. This, at any rate, is a position that Brooks has supported. Clark, however, does not and need not agree with this sort of representational skepticism (Clark and Toribio 1994).

https://doi.org/10.1086/652963 Published online by Cambridge University Press

4. Rowlands on the Barking Dog Principle. As I mentioned previously, the barking dog principle is more general than the 007 principle. In general, Rowlands thinks, *manipulationist strategies* (his term for strategies that put the environment to use) are more likely to evolve than internalist strategies, and, because a strategy like Reflex's is manipulationist whereas Cogito's is internalist, Reflex-style cognition is more likely to evolve than Cogito-style cognition.[7] Of course, because the 007 principle is a special case of the barking dog principle, reason to doubt the latter principle is also reason to doubt the former.

The distinction between manipulationist and internalist strategies that Rowlands draws is not black and white, and this becomes a source of trouble, but the intuition behind the distinction is clear enough. A strategy is a trait that has been selected to address a particular selection problem. Thus, although the term 'strategy' often conveys a kind of behavior, Rowlands intends a broader usage, to include morphological, anatomical, and other sorts of traits. Thus, on Rowlands's view, the dung fly's mating behavior, the cassowary's horn, and the thick musculature of the heart's ventricles are all strategies. A strategy is manipulationist if it exploits features of its environment in the course of addressing the problem it has been selected to solve. For instance, Rowlands claims that the beaver's dam-building strategy is manipulationist. The selection problems the beaver faces are the acquisition of food and avoidance of predators. Beavers accomplish both these goals by manipulating their environments. They fell trees, thus damming a stream, thus creating a pond, thus creating a safe haven while at the same time expanding the area of food-bearing shoreline.

Similarly, according to Rowlands, parasites often engage in manipulationist strategies. The cuckoo lays its eggs in a host's nest, thus foisting the task of child care on members of another species. The parasitic fluke *Dicrocoelium dendriticum* lodges itself into the brain of an ant, causing the ant to climb to the tops of grass stems, where it is more likely to be consumed by sheep, the fluke's ultimate host. Thus, the fluke quite literally manipulates the ant to accomplish its goal of inhabiting a sheep.

In contrast to manipulationist strategies are those that do not exploit features external to an organism in order to overcome a selection problem. Rowlands labels this second form of strategy *internalist*. Rowlands illustrates internalist strategies with counterfactual organisms, organisms that

7. To be fair, Rowlands cautions against the "Panglossian" assumption that because manipulationist strategies are more cost effective, they will inevitably evolve (1999, 96). Hence, his official line is that if an organism evolves the most efficient evolutionary strategy, then it will adopt a manipulationist strategy (101). However, Rowlands also thinks that, typically, organisms will evolve a manipulationist strategy (101).

https://doi.org/10.1086/652963 Published online by Cambridge University Press

are counterparts to those that, in the actual world, evolved manipulationist strategies. For instance, the *super beaver*, faced with the same tasks of finding/transporting food and avoiding predators as the actual beaver, took a different evolutionary path toward a solution. Whereas the actual beaver evolved large teeth, a flat tail, and the musculature necessary for effective use of these traits, the super beaver evolved a suite of traits that enabled it to eat and avoid predators exactly as successfully as the actual beaver but without having to rely on the manipulation of external structures. Rowlands describes the elements involved in the super beaver's internalist strategy:

> Firstly, the super beaver will require more powerful muscles for dragging its food on the long overland journeys it must make. Thus its limbs and torso must become more powerful (i.e. larger). Secondly, it must also possess the capacity to escape from the predators it will inevitably meet on these long overland journeys. Thus, it might have to become quicker, which might in turn generate the need for longer limbs, etc. It might also have to become more intelligent, thus creating a need for brain encephalization. . . . Encephalization entails a larger brain, which, in turn, requires stronger, i.e. larger, muscles, and so on. (1999, 76)

Rowlands asks us to imagine that actual beaver and super beaver are, as a result of their divergent evolutionary paths, equally successful in procuring food and avoiding predators, but the former uses a manipulationist strategy—altering its environment to ease the solution of a given selection problem—whereas the latter does it all on its own, that is, without relying on structures in its environment for assistance.

The actual fluke's internalist counterpart, because it shuns "outside" help, will also have evolved an elaborate suite of traits that are not necessary if relying on a manipulationist strategy. The super fluke's internalist strategy "would require investment in structures or mechanisms which enable the fluke not only to escape from the ant, but also to survive in the 'wild' until ingested by a suitable itinerant ungulate. It is fairly clear that this is a much more complex, hence more difficult, strategy to execute than the actual strategy employed by *D. dendriticum*" (1999, 86). In sum, actual fluke's strategy is manipulationist because it uses the ant to go where it wants to go. Super fluke's strategy is internalist because it goes where it wants to go under its own steam. The difference is between climbing a mountain on your own and allowing a ski lift to get you to the top.

I believe that Clark would be very sympathetic to the considerations in favor of manipulationist strategies that Rowlands raises. Super beaver plays Cogito to actual beaver's Reflex. Natural selection ought to produce

https://doi.org/10.1086/652963 Published online by Cambridge University Press

manipulators rather than internalizers, for just the same reason that engineers should build robots that use the world as its own best model rather than opting for more costly and computationally demanding strategies.

5. Problems with the Barking Dog Principle. Restating Barking Dog (BD) at this point will be helpful: "If it is necessary for an organism to be able to perform a given adaptive task T, then it is differentially selectively *dis*advantageous for that organism to develop internal mechanisms sufficient for the performance of T when it is possible for the organism to perform T by way of a combination of internal mechanisms and manipulation of the external environment" (Rowlands 1999, 80). In this section, I present three reasons to be suspicious of BD. As I mentioned earlier, because the 007 principle is just a special case of BD, applied more narrowly to cognition, doubts about BD must also create doubts about 007. The reason to question BD, and hence 007, emerges from reflection on counterfactuals implicit in BD. In effect, BD invites us to imagine a point in time in a species' evolution when members of the species confront a selection problem for which they have not yet evolved a solution.[8] If BD is true, we should expect that any solution to this problem that does evolve by natural selection will be manipulationist.

To fix ideas, let us suppose that we can look back in time to the moment when the beaver lineage is confronted with the selection problem that, according to Rowlands, led to the evolution of the actual beaver's manipulationist strategy. Already, we can see that there is something artificial in the way BD has us thinking about matters. Presumably the beaver's ancestor, protobeaver, already had strategies in place to collect food and avoid predators. If there are any universal selection problems, surely these two would be among them. Thus, protobeaver had some strategy by which it collected food and avoided predators. The point in time we are imagining is therefore not one when the ancestors of current beavers had no solution to the problems of collecting food and avoiding predators.

More sensibly, BD applies at the moment in time when members of the protobeaver population evolve either a manipulationist strategy or an internalist strategy, where the manipulationist strategy is that of the actual beaver and the internalist strategy is that of super beaver. BD predicts that over time actual beaver's manipulationist strategy will move to fixation in the population (barring issues such as heterosis, etc.). But why should this be true? Lurking in the bushes are a number of assumptions

8. I will assume that Rowlands is using 'develop' as a synonym for 'evolve'. Rowlands's statement of BD seems to suggest that organisms, rather than species, evolve. I will assume that although natural selection can operate on individual organisms, evolution takes place at a population level.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

to which Rowlands must be committed. Once exposed, BD loses much of its plausibility. The assumptions are these:

i) the ancestral condition does not bias the direction of evolution;
ii) environmental conditions do not bias the direction of evolution;
iii) selection problems remain fixed in the course of evolution.

In the remainder of this section, I will examine each of these assumptions.

5.1. An Assumption about the Ancestral Condition. We are supposing that protobeaver already has some strategy for acquiring food and avoiding predators. If protobeavers have survived this long, however long that may be, they must have evolved some strategies for nourishment and protection. These strategies mark the ancestral condition from which the strategies of actual beaver and super beaver emerge. Whatever the ancestral condition, it is likely to bias the direction of evolution in future generations. Thus, if protobeaver resembled a squirrel, which is already equipped with strong teeth and a tail, one might suppose that fewer steps would be involved in the evolution of actual beaver's strategy from this ancestral population than there would be in the evolution of super beaver's strategy. However, if the ancestral population resembled a jackal, the evolution of super beaver's strategy might well take fewer steps.

Of course, evolution does not always follow the most parsimonious path, that is, the path with fewest evolutionary steps. However, BD requires something quite implausible. It requires that regardless of the ancestral condition, evolution will always lead toward a manipulationist strategy. The important point is that even if, as Rowlands argues, manipulationist strategies are more energy efficient than internalist strategies, this is no guarantee that manipulationist strategies will evolve.[9] To insist otherwise is to endorse an extreme (and implausible) form of adaptationism—an adaptationism asserting that natural selection will always produce the fittest trait regardless of constraints that ancestral character states impose.

A response to this point might claim that because selection will always favor manipulationist strategies, protobeaver will already be using a manipulationist strategy, and thus actual beaver's strategy is more likely to move to fixation than is super beaver's. Moreover, protobeaver's strategy

9. As I noted in n. 7, Rowlands is aware of the dangers of unwarranted adaptationism. However, while voicing caution, he also seems prepared to accept that selection will typically produce manipulationist strategies. If attribution of the criticism above cannot be made to stick to Rowlands, it has a better chance of adhering to Clark, who appears to offer his 007 principle without qualification.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

is manipulationist because its ancestors had a manipulationist strategy, and so on.

There are several problems with this response, some of which will become clear in my discussion of the other assumptions on which BD depends. However, worth noting is that this response is peculiar in light of the considerations on which Rowlands relies to motivate BD. BD is supposed to be true because selection favors manipulationist strategies. The biological examples Rowlands discusses are all intended to highlight the idea that manipulationist strategies will evolve because they are fitter than internalist strategies. This is no doubt also behind Clark's confidence in 007. But the response above defends BD not on selectionist grounds but on grounds of ancestral constraints. On this proposal, the reason actual beavers evolved from protobeavers is not because the actual beaver strategy is superior to the super beaver strategy but because the protobeaver strategy biased evolution in the direction of the actual beaver strategy.

To summarize, BD assumes something unlikely: that the ancestral condition is neutral with respect to manipulationist and internalist strategies. When the ancestral condition is not neutral between these strategies, BD faces a dilemma. If the ancestral condition biases in the direction of a manipulationist strategy, BD is false because the manipulationist strategy does not evolve because (or, at least, only because) of a selection advantage. If the ancestral condition biases in the direction of an internalist strategy, BD requires that evolution take the less parsimonious route. Although evolution sometimes does take the less parsimonious route, this should happen less often than otherwise. Surely, evolution should not usually or always take the less parsimonious route when doing so involves a change from an internalist strategy to a manipulationist one, as BD requires.

5.2. An Assumption about the Environment. For sake of argument, let us suppose that the ancestral condition is neutral between manipulationist and internalist strategies. Protobeaver's strategy does not bias evolution in the direction of either actual beaver's or super beaver's strategies. I am not sure whether this is a real possibility given Rowlands's suggestion that strategies are either manipulationist or internalist, but perhaps matters are not so cut and dried. Or, perhaps protobeaver has a strategy that is one or the other but not so that a bias is created in the direction of the strategies of actual beaver and super beaver. If BD is true, actual beaver, but not super beaver, should evolve. But surely this depends on the nature of the environment in which protobeaver exists.

Suppose the streams in protobeaver's environment dry up during the summer months. In this event, manipulationist members of the population would be at a selective disadvantage compared to internalist members.

https://doi.org/10.1086/652963 Published online by Cambridge University Press

While actual beavers would starve and be subject to predation during the summer months, super beavers would be no less well off. They would continue to chop down trees, using their strong muscles guided by big brains to drag them to safety where they can then consume the tree's bark. The actual beaver's strategy would be at a selective disadvantage. BD would be false.

However, perhaps this objection is unfair. As stated, BD asserts that selection would prefer manipulationist strategies when it is possible to perform these strategies. In describing a case in which streams run dry in the summer, I have described a case in which a manipulationist strategy, or at least the one that actual beavers have adopted, is not possible.

But this response only invites harder questions. First, why should we expect that manipulationist strategies will often be available? Both BD and 007 suggest that manipulationist strategies will be the rule rather than the exception. This in turn assumes that nature, for some reason, has laid itself bare to manipulation. Perhaps this is true. Indeed, perhaps this explains why Earth is the kind of planet on which life has evolved. This is an interesting claim, but one for which I would like to see some evidence.

Second, how are we to understand BD's claim that manipulationist strategies will evolve when possible? Suppose it were possible to prevent streams from running dry in the summer. Perhaps a system of wells and pumps could keep the streams running year-round. It is then possible for actual beavers to maintain their manipulationist strategy throughout the year. Of course, in order to do so, actual beavers would have had to develop more than large teeth and flat tails. Presumably, they would have had to encephalize to a much greater extent than super beavers would have. Moreover, they would have had to develop motor skills sufficient for the construction of wells and pumps. If this is what is required for the adoption of a manipulationist strategy, my bet is that selection would have favored super beaver's internalist strategy.

In order to prevent objections like the preceding, BD requires modification. Rather than stating simply that manipulationist strategies will evolve when possible, it should say something like "manipulationist strategies will evolve when possible and when energetically less costly than internalist strategies." Or, because "energetically less costly" would seem to entail possibility, BD should stipulate that manipulationist strategies will be preferred to internalist strategies when they are energetically less costly. Clark's mention of cost in his statement of the 007 principle seems to endorse this idea as well.

But at this point BD loses its interest. Rowlands, following Van Valen (1976), conceives of fitness as energy maximization. Thus, BD, when modified to avoid the objection just raised, merely states that fitter strategies will be selected over less fit strategies. BD was intended as an empirical

https://doi.org/10.1086/652963 Published online by Cambridge University Press

claim about the relative advantage that manipulationist strategies have over internalist strategies, but, on further reflection, it becomes a truism. What would be of interest is the claim that manipulationist strategies are always more energy efficient than internalist strategies. But this is precisely the point that the discussion above calls into question: manipulationist strategies will not, on average, be selected over internalist strategies when they require for their evolution far greater resources than internalist strategies.

5.3. An Assumption about Selection Problems. Suppose again that some members of the protobeaver population evolve a manipulationist strategy and others an internalist strategy. Assume as well that the ancestral strategy by which the protobeaver collects its food and avoids predation is neutral between the manipulationist and internalist strategies in the way I described above. Will, as BD asserts, the manipulationist strategy be fitter than the internalist strategy?

Rowlands's affirmative answer to this question assumes that the selection problems facing both sorts of members of the population remain the same. That is, Rowlands assumes that both actual beavers and super beavers will continue to rely on the same food resources and will face the same kinds of predators. The manipulationist strategy is superior to the internalist strategy, given that the goal of each is to harvest trees and avoid terrestrial predators.

But why assume that strategies evolve against a fixed background of selection problems? Typically, as species evolve, the nature of the selection problems they face change as a result of their own evolution. This phenomenon is most conspicuous in evolutionary arms races. Species X preys on species Y and evolves a trait that makes predation on Y easier. Consequently, there is selection in the Y population for any trait that thwarts members of X. Selection will then favor traits in the X population that can overcome the defensive strategies that have most recently evolved in Y. The point is that because evolution occurs, new selection problems emerge.

This is true, not just in the case of arms races. One might imagine that selection for flight brought along with it selection for good vision. Eagles soar high enough above the ground so that their prey cannot see them. But this does the eagle no good if it cannot see its prey. Perhaps protoeagles did not fly far above the ground and so did not need sharp vision to find their prey. The problem of seeing more acutely arose in step with the strategy of flying to greater heights. As species evolve, so do the selection problems that their members face.

With this point in mind, Rowlands's assumption that members of a population that move toward internalist strategies will continue to face

https://doi.org/10.1086/652963 Published online by Cambridge University Press

the same selection problems that face their manipulationist siblings appears dubious. A more likely story is this. Members of the protobeaver population that have adopted the actual beaver strategy will cut trees and create dams in order to collect food and avoid terrestrial predators. Other members of the population, also descended from protobeavers but not necessarily super beavers, adopt an internalist strategy that changes the nature of the selection problems that these organisms must face. Whatever this internalist strategy, there is no reason to expect that it will be of much use in harvesting trees and avoiding terrestrial predators. Indeed, the strategy may be terribly suited for these tasks. But this does not mean that this internalist strategy is maladaptive. It may be better adapted for the collection of certain foods and the avoidance of certain predators than the actual beaver strategy is for the harvesting of trees and avoidance of terrestrial predators. Rowlands goes wrong in assuming that internalist and manipulationist strategies will face precisely the same selection problems (in particular, those better addressed by manipulationist strategies). But there is no reason to expect that members of a population that have made an internalist response to selection problems must continue to respond to the same selection problems that face their manipulationist siblings.

6. More on the Manipulationist versus Internalist Distinction. If the barking dog principle were correct, manipulationist strategies would be evolutionarily probable, at least in those cases in which natural selection is the primary force behind evolution. Similarly, the 007 principle predicts that cognition will more often be Reflex-like than Cogito-like. I have argued that evolutionary considerations on their own do not favor manipulationist strategies. Whether a manipulationist strategy evolves depends not only on whether natural selection is operative but also on (i) the ancestral character states on which natural selection is operating, (ii) the friendliness of the environment toward manipulationist strategies, and (iii) the changing nature of selection problems as a result of evolution. Supporting my arguments is nothing more than a realistic appreciation for how populations evolve under the control of natural selection.

Perhaps, however, problems with BD are visible from an even more accessible vantage. One need not be a scholar of evolutionary theory to take note of an obvious difficulty BD faces. If natural selection is the primary driver of evolution, as many believe, and if BD is true, why are there so many internalist strategies in use by organisms today? Beavers, cuckoos, and flukes are objects of fascination precisely because of their ability to manipulate their environments to their own benefit. But what of all the species that Rowlands does not mention? In thinking about this, one might be struck by the number of species that do not use manipu-

https://doi.org/10.1086/652963 Published online by Cambridge University Press

lationist strategies. Cheetahs chase down their prey rather than building corrals into which they may be herded; crabs open shellfish with their claws rather than banging them open with stones; sea turtles use their flippers to burrow holes in the sand to bury their eggs rather than digging the holes with broken coconut shells; most birds raise their offspring the way most other species do—by themselves. The idea that manipulationist strategies should not only be common across the animal (and plant?) kingdom but should be nearly inevitable seems to fly in the face of what we know of natural history.

At this point, the need for a more precise analysis of the distinction between manipulationist and internalist strategies becomes imperative. Perhaps Rowlands might try to interpret the examples of internalist strategies that I mentioned above as in fact examples of manipulationist strategies. If manipulationist strategies are simply those that make any use at all of features external to organisms, then indeed all the examples I just mentioned are manipulationist. The cheetah does not eat itself. It solves the problem of food acquisition by eating objects outside itself, and so does the crab. Birds and turtles lay their eggs in environments external to themselves. In fact, respiration might count as a manipulationist strategy insofar as it makes use of gases outside the body.

The distinction between manipulationist and internalist strategies doubtless has some intuitive appeal. The cuckoo, for example, appears to make clever use of the environment in ways that a more typical bird does not. But Rowlands needs to say more about what this cleverness consists in such that a cuckoo's parental care strategy is manipulationist but a typical bird's strategy of building a nest and laying eggs in it is not. Alternatively, if the typical bird's strategy is manipulationist, Rowlands needs to tell us why not every strategy, insofar as every strategy would seem to depend to some extent on the external world, is manipulationist.

7. Evolution and Minds. Evolutionary theory does make predictions about the trajectories of traits in a population. Island dwarfism, for instance, is a well-documented case. Large mammals tend to shrink when isolated in small environments. Presumably, this is because selection favors those mammals that require less food in an environment with strictly limited resources. Thus, elephants became small in Crete and Cyprus. However, the prediction that large mammals situated on small islands will shrink is unlike the prediction that the 007 and barking dog principles make. The prediction of dwarfism starts with an assumption about the ancestral condition, that is, that the mammals are large. It also starts with an assumption about the environment, that is, that resources are insufficient to provide for the number and size of mammals on the island. Finally, it starts with an assumption about the constancy of the selection

https://doi.org/10.1086/652963 Published online by Cambridge University Press

pressure weighing on the evolving population, that is, that food is scarce. These assumptions explain why smaller animals will have a competitive advantage over larger ones.

In contrast to the prediction of island dwarfism, the prediction that organisms will tend toward manipulationist strategies seems quite naked. Neither the 007 principle nor the barking dog principle situates itself within a body of assumptions that might compel acceptance. Whether organisms will evolve manipulationist strategies depends on which strategies are already present in the population, on whether the environment lends itself to manipulationist strategies, and on whether evolution will change the character of the selection pressures that face the population. Because ancestral traits might be internalist, and environments might not lend themselves to manipulation, and because selection pressures might change as populations evolve, the hypothesis that evolving populations will always or usually fix on manipulationist strategies is dubious. Similarly, the suggestion that nature will always or usually select manipulationist cognitive strategies fares no better.

As I mentioned in section 2, the 007 and barking dog principles have played a role in defenses of the extended cognition thesis. For Clark and Wilson, because the 007 principle is true, cognitive systems evolve to become embedded in environmental structure, which in turn leads them to become extended cognitive systems in cases in which the environment bears the right sort of functional integration with the rest of the system (perhaps one that involves an interactive loop). For Rowlands, the barking dog principle ensures the possibility of cognitive systems that are capable of manipulating those pieces of the environment that carry task-relevant information. When an organism uses a manipulationist strategy to complete a cognitive task, "these processes are not located purely inside the skin of the cognizing organism; they straddle both internal and external components" (1999, 100).

What becomes of the thesis of extended cognition if, as I have argued, evolutionary processes do not inevitably march cognitive systems toward manipulationist strategies? Clearly, this criticism does not by itself show the thesis to be in trouble. Although we now have grounds to doubt the conditional statement that "if a cognitive system has evolved by natural selection, it will employ a manipulationist strategy," nevertheless, as a matter of fact, evolved cognitive systems may employ manipulationist strategies. If, by chance, exploitative cognitive systems did evolve, then defenders of extended cognition might well rest content—extension begins with embeddedness, and what difference does it make where embeddedness comes from?

As a response to some of the criticisms I made of 007 and the barking dog, this position is reasonable. At best, rejection of these principles would

https://doi.org/10.1086/652963 Published online by Cambridge University Press

weaken the grounds for making predictions about the frequency of extended cognition in a given population. If natural selection does not, as a general rule, favor manipulationist strategies, then the fact of evolution is no motivation for expecting extended cognition systems to be common. But, this point need not trouble those who wish to make claims about actual cognitive systems, rather than cognitive systems that have not evolved.

Among psychologists and philosophers there is a sense that an evolutionary backing lends "street cred" to theories of the nature of mind. For instance, some who defend a modular theory of mind—according to which the mind consists in hundreds or thousands of simple, relatively isolated, and domain-specific processors—find support from within evolutionary theory (see especially the work of the evolutionary psychologists). As in the present case concerning extended cognition, the ties between modularity and evolution are not as obvious as they are given to be and often depend on historical claims that extend beyond the reach of current evidence or conceptual claims that require further refining.[10]

One lesson to draw from the present discussion is that the relationship between a theory of mind and the theory of evolution is complicated. Whereas a theory of mind that is inconsistent with evolutionary theory may be easy to spot, a theory of mind that actually follows from the theory of evolution is much harder to spot. Evolution could have taken minds in many different directions. Principles like 007 and the barking dog would impose blinders on evolution, preventing it from developing cognitive strategies that would in fact be fitter than manipulationist strategies due to phylogenetic constraints, uncooperative environments, or shifting selection problems. The number of factors that natural selection must "consider" when creating a mind must be staggering. It is hard to believe that evidence will ever surface to show a likely route from evolution to embedded minds (or modularity, for that matter).

But there is a more hopeful lesson for the thesis of extended cognition. In the end, the thesis does not need the support of evolutionary theory. The main arguments for extended cognition hinge on considerations I discussed in section 2, having to do with the nature of the connection between an organism and its environment and the meaning of the claim that cognition extends. Absent from this list of considerations is whether natural selection would favor manipulationist strategies. Quite simply, the case for extended cognition is independent of whether natural selection should tend to favor manipulative strategies. The fact is that many cog-

10. For critical discussion of these issues, see Shapiro and Epstein (1998), Buller (2005), and Richardson (2007).

https://doi.org/10.1086/652963 Published online by Cambridge University Press

nitive systems do exploit the environment. This, not evolution, should be the starting point for discussion of extended cognition.

REFERENCES

Adams, Fred, and Kenneth Aizawa. 2001. "The Bounds of Cognition." *Philosophical Psychology* 14:43–64.

———. 2008. *The Bounds of Cognition*. Malden, MA: Blackwell.

———. 2010. "Defending the Bounds of Cognition." In *The Extended Mind*, ed. Richard Menary. Cambridge, MA: MIT Press.

Block, Ned. 2005. "Review of Alva Noë." Review of *Action in Perception*, by Alva Noë. *Journal of Philosophy* 102:259–72.

Brooks, Rodney. 1991. "Intelligence without Representation." *Artificial Intelligence* 47:139–59.

Buller, David. 2005. *Adapting Minds: Evolutionary Psychology and the Persistent Quest for Human Nature*. Cambridge, MA: MIT Press.

Clark, Andy. 1989. *Microcognition: Philosophy, Cognitive Science and Parallel Distributed Processing*. Cambridge, MA: MIT Press.

———. 1997. *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.

———. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.

Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58:7–19.

Clark, Andy, and Josefa Toribio. 1994. "Doing without Representing?" *Synthese* 101:401–31.

Goldin-Meadow, Susan. 2003. *Hearing Gesture: How Our Hands Help Us Think*. Cambridge, MA: Harvard University Press.

Menary, Richard. 2007. *Cognitive Integration: Mind and Cognition Unbounded*. New York: Palgrave Macmillan.

Richardson, Robert. 2007. *Evolutionary Psychology as Maladapted Psychology*. Cambridge, MA: MIT Press.

Rowlands, Mark. 1999. *The Body in Mind: Understanding Cognitive Processes*. Cambridge: Cambridge University Press.

Rupert, Robert. 2004. "Challenges to the Hypothesis of Extended Cognition." *Journal of Philosophy* 101:1–40.

Shapiro, Lawrence. 2010. *Embodied Cognition*. London: Routledge.

Shapiro, Lawrence, and William Epstein. 1998. "Evolutionary Theory Meets Cognitive Psychology: A More Selective Perspective." *Mind and Language* 13:171–94.

Shapiro, Lawrence, and Shannon Spaulding. 2009. Review of *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, by Andy Clark. *Notre Dame Philosophical Reviews*, http://ndpr.nd.edu/review.cfm?id=16245.

Van Valen, Leigh. 1976. "Energy and Evolution." *Evolutionary Theory* 1:179–229.

Wilson, Robert. 2004. *Boundaries of the Mind: The Individual in the Fragile Sciences; Cognition*. Cambridge: Cambridge University Press.

Wilson, Robert, and Andy Clark. 2009. "How to Situate Cognition: Letting Nature Take Its Course." In *The Cambridge Handbook of Situated Cognition*, ed. Philip Robbins and Murat Aydede, 55–77. Cambridge: Cambridge University Press.

https://doi.org/10.1086/652963 Published online by Cambridge University Press