

teorema

Vol. XXXVII/2, 2018, pp. 187-202

ISSN 0210-1602

[BIBLID 0210-1602 (2018) 37:2; pp. 187-202

Affording Affordances

David Spurrett

RESUMEN

Un rasgo sorprendente de la última versión del ‘gran cuadro’ proporcionado por Dennett de la evolución de la vida y de la mente es la frecuente referencia a los ‘posibilitadores’ (*affordances*). Un posibilitador es, hablando de forma aproximada, una posibilidad de acción para una criatura en un entorno. Dada la existencia de más de una posibilidad de acción, una buena pregunta es: ¿qué hará efectivamente la criatura? Argumento que los posibilitadores plantean un problema de selección y que una buena solución a este problema del diseño de la mente es implementar un sistema de preferencias.

PALABRAS CLAVE: *posibilitadores, selección de la acción, preferencias, caminos comunes finales.*

ABSTRACT

A striking feature of the latest version of Dennett’s ‘big picture’ of the evolution of life and mind is frequent reference to ‘affordances’. An affordance is, roughly, a possibility for action for a creature in an environment. Given more than one possibility for action, a good question is: what will the creature actually do? I argue that affordances pose a problem of selection, and that a good general solution to this problem of mind-design is to implement a system of preferences.

KEYWORDS: *Affordances, Action Selection, Preferences, Final Common Paths.*

Daniel Dennett’s new book, *From Bacteria to Bach and Back* [Dennett (2017)], which I’ll sometimes refer to simply as ‘*BBB*’, is one of his ‘big unified picture’ books, like *Consciousness Explained* [(1991)] or — even more so — like *Darwin’s Dangerous Idea* [(1995)]. While there are many new details and modifications in *BBB*, the broad outline is familiar because it is the view, or at least major parts of it, that he’s been defending for many years. Darwin and Turing are still central heroes, dualists and supernaturalists remain key targets of an ambitious and wide-ranging selectionist naturalism about life, mind and language.

A key reason that this book exists then, is not so much that Dennett has a big new thing to say, but that he’s taken the lessons of many

years of attempts at explanation and persuasion, with their attendant successes and failures, to work out substantially updated expositions of his views. He says as much in the opening pages. For the reader already acquainted with the earlier work, perhaps especially for the one in broad agreement (which includes me) the pleasure of this book is akin to that of reading a new translation of a familiar work, or seeing a fresh adaptation of a favourite play.

BBB isn't merely a retread, though: there are many changes of detail. Some are subtractions. So there's hardly any mention of 'figment' or the 'library of Babel', none of 'universal acid', the sphex wasp, or of Orwell and Stalin. There are also differences of emphasis, including an expanded and updated defense of the meme-concept, and an extended discussion of what is distinctive about encultured and tool-wielding human minds. A long familiar theme warning against supposing comprehension where uncomprehending competence is sufficient is forcefully emphasised. One change that I found especially striking is the inclusion and central use of a notion of 'affordance' in Dennett's recent exposition. There's no explicit mention of affordances, or of J. J. Gibson who coined the term, either in *Consciousness Explained*, or in *Darwin's Dangerous Idea*. Most of what I have to say here will focus on the place of affordances in Dennett's argument. Affordances are generally understood as a kind of *possibility* for an organism (given its needs and capacities) in a specific environment. They're often, but not always, possibilities for some kind of activity or behaviour, and a common homely example is that a chair affords sitting for a person of the right size.

In what follows I spend a little time getting clearer on what affordances are supposed to be, especially for Dennett. Then I argue that thinking about cognition in terms of affordances requires facing up to the problem of selection, whether between simultaneous and incompatible 'active' affordances, or of how much and when to perform the behaviour associated with a single one, or several that aren't mutually exclusive. I go on to suggest that a system of *preferences* (sub-personal representations of value) is a good general design response to the selection problem, and that preferences can, furthermore, be a satisfying example of competence without comprehension. I regard the view I'll suggest about preferences as Dennettian, and it draws direct inspiration from some of his earlier work. Indeed, I struggle to see how the notion of affordances can do the required explanatory work *without* something like preferences, so I think it is worth making the reasons for this explicit.

I. AFFORDANCES

The English verb ‘afford’ has two related primary meanings. In one, to ‘afford’ something is to *have the means* for it, most commonly with reference to a budget of money or time. (For example ‘I can’t afford to sit and chat, I have a train to catch’.) In the other, to ‘afford’ something is to *provide or supply* it. (As in, ‘This ledge affords a view of the whole valley’.) Gibson mostly used the latter sense, and also coined the noun ‘affordance’ to refer to the provision of some possibility to a creature given its capabilities and priorities [Gibson (1979)]. Since Gibson introduced the term, it has been an important concept in ecological psychology, and enjoyed significant attention in discussions of distributed, situated or embodied cognition [e.g. Clark (1997)], as well as by enactivists [e.g. Chemero (2009)]. Although the term has been in use for around four decades, there’s been a striking recent rise in activity, and it has featured in several high profile synoptic works on cognition, including Andy Clark’s *Surfing Uncertainty* (2016), Ruth Millikan’s *Beyond Concepts* (2017) and Daniel Dennett’s *From Bacteria to Bach and Back* (2017). It has also gained modest currency in parts of cognitive neuroscience [e.g. Cisek (2007)]. None of these works is ‘strictly’ Gibsonian, and all put the notion of an affordance to work in distinctive or qualified ways. A little orientation will help focus on what Dennett specifically is up to.

Gibson’s notion of an affordance can be given engagingly simple glosses, such as that the “affordances of the environment are what it offers animals, what it provides or furnishes, for good or ill” [Gibson (1979, p. 127)]. But these need to be hedged or qualified, to guard against tempting interpretations that Gibson was keen to reject. Gibson was a kind of direct realist about perception, and his general psychological views were in some respects anti-behaviourist and (later) anti-cognitivist, especially in rejecting representational conceptions of cognition. So, for example, he took pains to reject the view that affordances were *stimuli* [e.g. (1979), p. 140]. An affordance isn’t just a property of the environment, nor is it just a psychological (or psycho-mechanical) property of an organism. As Gibson puts it, an affordance “is neither an objective property nor a subjective property; or it is both if you like. An affordance cuts across the dichotomy of subjective–objective and helps us to understand its inadequacy. It is equally a fact of the environment and a fact of behavior” [Gibson (1979), p. 129].

Gibson counted a very wide range of relations and possibilities as instances of affordances, including the ground affording standing, the air affording breathing, surfaces affording *falling off*, other organisms affording being eaten, as well as affordances of social life and tool use. Not all affordances are positive or beneficial, and it is part of Gibson's own view that the *value* of an affordance is directly perceived [Gibson (1979), p. 137]. As Gibson speaks of them, affordances specify things a creature might do and also things that might happen to it (in a particular environment, given its body) that it in some sense *appreciates*. These can be good, like breathing and mating, or bad, like falling down or being eaten, and the valence is an aspect of what is perceived.

The recent appropriations of affordance-talk that I'm interested in here, especially by Clark and Dennett, acknowledge Gibson but are careful to signal that they're not 'strict' Gibsonians. Both Clark and Dennett are kinds of cognitivist, even though neither is friendly to the sorts of 'Classical Sandwich' conceptions of cognition rejected by Gibson. Dennett has, furthermore, long-standing if qualified sympathies with behaviourism. In his earlier *Being There*, Clark credits Gibsonian affordances as being a 'direct inspiration' for his notion of an 'action-oriented representation' [Clark (1997), p. 172]. That type of representation, recall, is not a passive world-model requiring interpretation in advance of action-specification. Rather, it is a structure including both aspects of the transduced environment and possible actions in a form that makes the representation itself fit to serve as a controller [Clark (1997), pp. 47-51]. His current contention is that prediction-driven learning is well-suited to giving "a grip" on affordances, still understood as organism-relevant "opportunities for action and intervention" [Clark (1997), p. 171, xv]. Clark, unlike Dennett, also refers to the work of Paul Cisek, who has offered a view of the neural implementation of action selection called the 'affordance competition hypothesis' [Cisek (2012)].

Dennett's deployment of the notion of affordances is unsurprisingly, given their common ground, not far from Clark's. Both embrace the *one* key Gibsonian commitment that affordances are sensory-motor relations, specific to combinations of environments and particular organisms, while avoiding the direct realism and anti-representationalism. Where Clark notes Gibson as an inspiration for his notion of an 'action-centric representation', Dennett says that his own related notion of "useful information" has Gibsonian antecedents [Dennett (2017), p. 180]. For his part Dennett glosses affordances as "the relevant opportunities in the environment of any organism: things to eat or mate with, openings

to walk through or look out of, holes to hide in, things to stand on, and so forth” [Dennett (2017), p. 126]. And tracking affordances, understood as organism-specific relations to useful information both positive and negative, is a key function, perhaps *the* key function, of brains: “Brains are control centers for dealing swiftly and appropriately with the opportunities and risks — the affordances — of a mobile life” [Dennett (2017), p. 227]. Dennett, like Gibson, contemplates a wide-ranging and inclusive notion of affordances, even including affordances associated with appreciating jazz improvisation.

In Clark and Dennett the accounts and glosses of affordances focus largely on affordances as *provision* or possibility. That’s well and good if you’re mostly trying to explain what an affordance is, and how to do without the notion of a brute environment that organisms cope with by making passive models that then stand in need of further interpretation before they can make any contribution to action selection or execution. But that can’t be the full story about affordances and the evolution of minds, because of what I’ll call the problem of affordance selection.

II. SCARCE MEANS WITH ALTERNATIVE USES

Affordance-talk is, as we’ve seen, almost always about *possibility*. Gibson spoke of affordances as what the environment ‘offered’, leaving open what, if anything, was ‘taken’ by the organism. Clark and Dennett alike refer to them in terms of ‘possibilities’ and ‘opportunities’. An obvious question, then, is given one or more affordances, what will the creature *actually* do? I’ll call this the problem of affordance selection.

To begin thinking about this, suppose there’s only one affordance. This might be because the creature only has one relevant capacity or behaviour, or because the world is so disposed as to provide only one affordance to that creature at some time. Even in this minimal case, there’s a question about *whether* the behaviour will be produced, and if it is, when it will, how intensely, and for how much of the time that the affordance is ‘active’. (To the experienced lab pigeon, a key affords pecking, but rewarded keys don’t generally get pecked *continually*.) Even in the single-affordance case, it is clear enough that the questions about what will be done, when, and how much, are partly *economic*. Producing behaviour usually draws on (even if its consequences might also replenish) a finite budget of energy, water or something else, and use up all-too-finite time.

It won't *always* be the case that an animal can afford to produce the behaviour that is afforded.

The questions become more interesting if we consider cases — arguably more common for most creatures most of the time — where *multiple* affordances are in play. While some combinations of behaviours can be produced together (like walking and chewing gum)¹ the economic dimension is most striking for mutually exclusive behaviours (like fighting and fleeing). The mutual exclusivity of two behaviours might arise in more than one way, including because different occurrent features of the world simultaneously afford mutually exclusive activities, or because one or more of them affords multiple candidate activities that are mutually exclusive. The hungry and thirsty creature that detects food to one side of its head and water to the other one is subject to two simultaneously incompatible affordances (if it has only one mouth). So too is the thirsty, drinking creature that detects a predator, and can either continue drinking or flee. What makes behaviours mutually exclusive in any given case could be a fact about the structure of the environment or of the creature's body, or a mixture. (The factors can also change. Standing in a queue and sending an email used to be mutually exclusive.) An important source of these constraints for some behaviours is the specific interdependencies and relations between the various degrees of freedom in an organism's own (extended) body. Carrying something in your mouth rules out making some vocalizations. Running at full speed requires co-ordination of both arms and legs, etc. None of this is controversial or novel, but it is worth emphasising that the mutual exclusivity of affordances — which just are possibilities for action given properties of both environment and organism — is partly contingent of the structure of the capacities of the organism *and* the structure of the environment.

One generic way that a creature can be more effective is by more often doing 'better' at producing appropriate behaviour given multiple affordances. (That is, of course, ultimately better in currency recognised in the game of life. I will take this as given.) Creatures that tend to favour drinking over eating when more dehydrated, or fleeing at the expense of water when predators are dangerously close, for example, can be expected to do better than ones that make less advantageous selections. The problem here is economic in a quite strict way - it concerns the allocation of scarce means (energy, time, bodily degrees of freedom) with alternative uses. And the allocations don't only have varying costs, they have varying returns (in rehydration, nutrition, staying alive, copulating,

etc.). Selection between affordances that is sensitive to the costs and returns is selection that tracks what the creature can afford.

Gibson, while more interested in perception than explaining cognition or behaviour, has a key *part* of the answer to the problem of selection, because of his view, noted above, that the *value* of an affordance, whether positive or negative, is itself directly perceived. (I say ‘part’ of an answer, because we’d still need to say something about what gets done with the appreciated values, including how exploitation and exploration are traded-off where that is a factor.) Dennett, as we’ve seen, does say that affordances can be positive or negative. But he doesn’t say explicitly that this fact about them is a component or aspect of the affordance considered as a cognitive process in the organism. For this reason, the picture we are given in *BBB* is a little cryptic about the problem of affordance selection.

That said, few philosophers have taken *selection* as consistently seriously as Dennett, who as I noted in opening has long argued that Darwinian processes of trial and error, variation and selection, at different scales (including time-scales) and in varying organisational hierarchies have extraordinary explanatory range. These processes include ‘vanilla’ selection on *Darwinian* agents whose behavioural dispositions are fixed by between-generation selection, and the ‘post-natal design-fixing’ by which *Skinnerian* agents tune their dispositions in light of a history of reinforcement, as well as the planning and selection processes in more sophisticated agents capable of various forms of ‘off-line’ learning, and exploitation of (and also by!) culturally transmitted replicators.

The distinction between *Darwinian* and *Skinnerian* agents goes back to his essay “Why the Law of Affect Will Not Go Away” [Dennett (1975)] which I was recently very pleased to find Peter Godfrey-Smith counted among his ‘favourite’ papers by Dennett [Godfrey-Smith (2018)]. That distinction was elaborated in *Darwin’s Dangerous Idea* (1995) and *Kinds of Minds* (1996) into what Dennett called the ‘Tower of Generate and Test’, and Godfrey-Smith goes on to argue, drawing on unpublished work by Russell Gray and Alex Taylor, in favour of a richer classification, arranged in a tree rather than a tower. There’s much of value and interest in his proposal, as well as Dennett’s reply [Dennett (2018)] but when I return to this topic towards the end of III, I’m going to focus on *Darwinian* and *Skinnerian* agents as Dennett describes them.

Both types (unless their capacities are hugely limited or their worlds uncommonly simple) face the problem of selection — between afforded behaviours — that I described above. I want to sketch a view, inspired

by the earlier work I've just drawn attention to, that I hope Dennett would find agreeable, and then briefly relate it to a concern that is in the foreground right through *BBB*, that we don't suppose comprehension where competence will do. The key idea is that what I'll call 'preferences', understood as cognitive states valuing possible behaviours,² or representing them as better or worse, can be an elegant solution to the problem of affordance selection.

III. AFFORDING PREFERENCES

Giant sea slugs (*Pleurobranchaea*) are carnivorous and will generally eat any animal matter they encounter, "including other sea-slugs and their eggs" [Manning and Dawkins (1998), p. 226]. An important exception to this indiscriminate behaviour is that they do not eat their own eggs during egg laying. This restraint is clearly adaptive: the genes of creatures that routinely consume their own descendants will tend to be poorly represented in subsequent generations compared to ones that don't. The mechanism stopping sea slugs from eating their own eggs, though, is a quite simple and direct inhibition. It operates without representations of the disutility of infanticide: during egg-laying their bodies release a hormone that inhibits movement of the mouth [Davis *et al* (1977)]. They don't eat *at all*, that is, while laying eggs. But there's no sense in which the return on eating their own eggs is represented as negative. Rather the affordance of the eggs right under them is simply turned off, by an inhibitory link strongly reminiscent of the representation-free 'subsumption' relationships once championed by Rodney Brooks [e.g. (1991)].³

The organism that solves behaviour allocation problems by methods such as these is prone to 'leaving money on the table', as they say to refer to a deal that doesn't extract some of the available returns. In this case, for example, the sea slug that happens to lay eggs while close to something nutritious that *isn't* its own eggs misses that tasty morsel and the calories it could have accumulated. Many creatures can, furthermore, demonstrably do better. In a series of important experiments Peter Shizgal and Kent Conover offered rat subjects mutually exclusive choices between trains of brain stimulation reward (BSR) pulses of varying lengths and infusions of sugar solution of varying sizes, as well as choices where one option returned a bundle reward including both BSR and sugar solution. They found that the observed choices, including cases where

one option was a bundle, were sensitive to opportunity cost (from the forgone lever), which is to say that the rats were capable of trading off the varying multi-modal returns quite finely and switching lever accordingly. As they reasoned:

In natural settings, the goals competing for behavior are complex, multi-dimensional objects and outcomes. Yet, for orderly choice to be possible, the utility of all competing resources must be represented on a single, common dimension [Shizgal & Conover (1996)].

What Shizgal and Conover mean by ‘orderly choice’ is simply choice that efficiently trades off costs and gains for rewards and punishments that differ in both magnitude and type. Orderly choice, that is, is efficiently solving the affordance-selection problem, if we regard the presented levers before the familiarised rat as affording pressing. I’m using the term ‘preferences’ to refer to the representations of utility that they have in mind. A preference is a real cognitive state, that somehow represents, encodes or tracks the (expected) value of local world-states or outcomes that an agent can detect or anticipate, or actions that it can perform. Preferences are, more or less, utilities.

Shizgal and Conover say that the value of competing resources “must” be represented on a single scale, suggesting that abstract utility representations are the *only* way to accomplish efficient allocations. Whether or not that is the case isn’t my concern here.⁴ For current purposes a weaker claim, that relatively abstract value-tracking — by means of preferences — *can* enable efficient selections, is sufficient. Before saying more directly about preferences, though, I need to spend a little time talking about final common paths.

Sherrington (1906) is generally recognised as among the first to articulate the key ideas here, independently of the topic of preferences. Consider a complex bit of animal anatomy with multiple degrees of freedom such as a primate forelimb with shoulder, elbow and a set of digits. Some combinations of allocations of the effectors, such as simultaneously flexing and relaxing the same muscle, or flexing those that would move a segment in one direction around a joint while not relaxing those that would move it in the opposite direction, are at least incompatible and could even be harmful. Sherrington argued that upstream of the specific effectors (in what he called the ‘afferent arc’) competing allocations (‘reflexes’) converged in a ‘final common path’, where only one reflex (as

opposed to some combination or sum) would gain control of the effectors downstream [e.g. Sherrington (1906), pp. 117-118].

The scope of the ‘common’ part of the expression ‘final common path’ depends on how many effectors or degrees of freedom have possibly incompatible uses given the candidate actions. This can easily be quite large, since some actions, such as reaching and grasping, require that multiple degrees of freedom be more or less simultaneously co-ordinated, as well as having others ‘co-operate’ by not interfering. (The grasping elements of a movement to pick a ripe piece of fruit won’t succeed if the rest of the arm is co-opted for ‘waving hello’, or if the legs move the whole body away.) A final common path is the form of a *possible* solution to these more general control problems too. It is a functional region of a control system in which events that prevent incompatible allocations of degrees of freedom, and which co-ordinate multiple degrees of freedom, can occur. While this line of reasoning could be stated and developed in a relatively abstract and anatomy-independent way, the historical version arose in the course of enquiry into the functions of specific parts of the spinal nerves and primary motor cortex.

We should be very wary of taking this reasoning to establish the need for *strict* and *total* final common paths in organisms with complex bodies. That would be like supposing a determinate ‘finish line’ for the control of a system spread out in space. The same considerations that weigh against assuming such a finish line for consciousness [Dennett (1991)] apply by analogy here. There will have to be some blurring and smearing of boundaries, and indeterminacy in some temporal ordering relations. That said, there are undoubtedly *local* neural final common paths, such as systems for controlling saccadic eye movements in creatures capable of them [Glimcher (2003)]. There are also distributed and hierarchically organised *approximate* ones in the somatotopic neural maps found in many real brains.

Final common paths enable, it is worth noting, an exemplary kind of competence without comprehension. Whether innate or acquired from experience, the fluent co-ordination of degrees of freedom to produce behaviours out of complex bodies, with only the rarest attempts at structurally conflicted activity — at least in experienced creatures — is remarkable. While it may be tempting to suppose that this depends on knowledgeable planning, the mundane truth is that brains and peripheral nervous systems are typically set up so that many inappropriate allocations are simply ruled out. In addition, and of special interest for my purposes here, final common paths provide a natural architectural or

functional ‘place’ for the operation of preferences. If there’s somewhere that mechanically incompatible actions are filtered out, then that somewhere might allow competition between options that are not ruled out to be expressed.

Some key early experiments in neuroeconomics depended on this very line of thinking to generate empirical hypotheses or to interpret their results [e.g. Platt & Glimcher (1999), Dorris & Glimcher (2004)]. In them monkey subjects were trained to express choices by making saccades to different targets. Saccades are a convenient behaviour precisely because they are controlled by small networks of muscles with their own series of neural topographic maps, corresponding to a two-dimensional ‘dart board’ of possible fixation targets relative to the skull. These maps are, in part, a final common path that generally prevents the eyes from attempting conflicted movements like simultaneously turning up and down. They’re also a kind of bottleneck for competition between fixation targets. (Do I continue looking at my conversational partner, or do I glance at the television screen off to one side?) It turns out that in the interval leading up to choice, the levels of activity in regions of the topographic maps corresponding to the saccadic targets under study varied with the expected return on that movement. Platt and Glimcher’s monkeys were ‘paid’ in juice, while Dorris and Glimcher’s subjects played an inspection game with returns in water. In later work Klein, Deaner and Platt (2008) found that single-unit activity in neurons specialised for saccades consistently reflected values of both social *and* fluid rewards, so this isn’t just something about fluid rewards.

This isn’t, furthermore, just a discovery about monkeys, and it isn’t just a discovery about saccadic movements. In neuroeconomic study of human choices, correlates of utility are usually sought further ‘upstream’ of anatomically detailed final common paths, because electrode recordings are rarely used in human neuroeconomic experiments, and choices expressed by button presses and the like don’t correspond to somatotopic maps that are as conveniently tractable as those for eye movements [Glimcher (2011), Chapter 7]. Even so, the evidence is rather compelling. Levy and Glimcher (2012) survey relevant experimental work up to 2012. First they detail studies showing that activity in the ventral striatum was positively associated with, among other things, monetary gains and losses, cumulative monetary rewards, anticipation of varying monetary rewards, expected values of uncertain monetary rewards, and discounted value of delayed monetary rewards. Second, they consider studies with at least one incentive other than money, including consumer goods, gusta-

tory rewards (water, juice, food), physical pain, social reputation, again finding consistent correlations. Whether or not Shizgal and Conover were correct to say that orderly choice means that there “must” be value representations in a common scale, the work of some neuroeconomists suggests that there is in fact one for a wide range of choice types. There is far less comparative neuroeconomics than I for one would like to read, but approximately analogous results have been found in other vertebrates and some invertebrates with indications that the neural systems across taxa share structural similarities, and similar functional roles for dopamine or related molecules. (This suggests that preference sensitivity, where it occurs, is an elaboration of older systems modulating motor control, rather than that everything that can move has preferences.)

I mentioned Paul Cisek in I above. His ‘affordance competition hypothesis’ [Cisek (2007)] is in the end a neuroscientist’s proposal for a way to make sense of the sort of neuroeconomic data I’ve described a fraction of, and other findings, *without* supposing some kind of ‘Classical Sandwich’ model of cognition. Cisek proposes, instead, that “the processes of action selection and specification occur simultaneously” [(2007), p. 1586] and argues that incoming sensory information selectively informs the generation of a number of incompletely specified behaviours, which may be released into execution prior to full specification. (That is, some of the specification would be handled *after* the behaviour was ‘launched’.) Reading his account of parallel, incomplete, competing behaviour specification processes with Dennett in mind, it is very tempting to call it a ‘multiple drafts’ model of affordance competition: “From this perspective, behaviour is viewed as a constant competition between internal representations of the potential actions which Gibson (1979) termed ‘affordances’” [Cisek (2007), p. 1586]. Cisek gets to affordances starting with neuroeconomic findings in clear view, and looking for a way of making sense of the brain without Classical Sandwiches. If we start with affordances, and face up to the problem of allocation, I think we should end up in approximately the same place as Cisek does. That is, I think that the argument of Dennett’s *BBB* requires that something like preferences be at work in *many* motile organisms.

Let us return to Dennett’s Tower of Generate and Test. It seems clear enough that what he calls *Skinnerian* agents have to have preferences. Recall that, unlike *Darwinian* agents all of whose design precedes the life of the agent, the more sophisticated *Skinnerian* agents undergo some ‘post-natal design-fixing’ in light of a history of reinforcement. In them, an additional form of selection is at work, based on assessment of

the consequences of their own behaviours. This selection requires an evaluation of the state that the agent is in, perhaps indexed to a recent salient behaviour of its own, to say whether that state is *better* or *worse* than some reference state. The reference state might be an ‘expected’ reward, and the comparison with the actual reward provide the basis for tuning the behavioural dispositions. It makes no difference to this if, in line with the way things are put in *BBB* we understand Skinnerian agents as affordance-tracking, and affordance-learning, machines. (Although an orthodox Gibsonian might be aghast at talk of *Skinnerian* affordance tracking.) ‘Expected’ here *doesn’t*, of course, mean a whole agent state of aware anticipation, merely some kind of sub-personal tracking. As Dennett says, the Skinnerian agent “doesn’t have to understand why it now prefers these tried-and-true behaviors when it does; it is the beneficiary of this simple design-improvement ratchet, its own portable Darwinian selection process” [Dennett (2017), p. 155]. Preferences, that is, can also deliver competence without comprehension: they’re represented, but not represented ‘to’ or ‘for’ the agent as a whole, merely serving traffic between the disposition-tuning systems. If we understand *Skinnerian* agents as in the affordance-tracking business, though, we’re committed to thinking of the affordances as being somehow evaluated.

It’s an interesting question whether *Darwinian* agents could occur with preferences. The fact that they don’t learn doesn’t by itself settle this question, because a Darwinian agent with anything but a trivially tiny behavioural repertoire or unusually simple environment has the problem of affordance selection, and the costs of and returns on its behaviours can be multi-modal and varied. Preferences that were not used to drive post-natal design fixing could still contribute to making the outcome affordance competition more advantageous in such agents, although I just don’t know if any real natural agents have this property. Reinforcement learning by ‘evolutionary methods’ searches the space of ‘policies’, which map states to actions, without the agents doing any learning themselves. Some ways of encoding policies *could* involve preferences, although in artificial systems policies can also be things like lookup tables. Perhaps *Darwinian* agents with proto-preferences account for some of the routes in design space that lead to *Skinnerian* agents. The indications that reward systems are elaborations of older motor control modulation systems encourage this view.

IV. CONCLUSION

I like to think that the picture I've sketched here is at least consistent with the much larger view of the evolution of life and mind laid out in *BBB*. As I hope I've made clear, the considerations that I've emphasised are directly inspired both by Dennett's own deep and wide-ranging interest in the importance of selection, and by specific arguments and thinking tools found in his own earlier work, perhaps especially including the 'Tower of Generate and Test', parts of which are incorporated into *BBB*. It at least isn't obvious that Dennett sees things this way, because outside the brief discussion of Skinnerian agents (which isn't framed in terms of affordances) his own discussion of affordances says relatively little about the problem of affordance selection, and the book as a whole barely mentions preferences (in the relevant sense) or cognate notions like utility.

I've sketched arguments that there is a problem of affordance selection, and that preferences are one way of solving that problem, which seems to have been taken — to varying extents — in many natural motile organisms. I think that naturalists who find affordance talk useful should both accept that there is a selection problem, and that preferences can help deal with it, despite the fact that in many cases people who talk about affordances do so without explicit acknowledgment of the problem of selection, or of preferences. As I noted early on, the word 'afford' can refer both to something *provided* or *supplied*, and to whether or not the *means* for some purpose are available. Selection is important partly in proportion to the severity of scarcity. Sensitivity to scarcity is, correspondingly, a potentially useful thing to have.⁵ So preferences can help an agent tell which affordances it can afford.

School of Philosophy & Ethics
University of KwaZulu-Natal
Mazisi Kumene Road
Durban 4041, South Africa
E-mail: spurrett@ukzn.ac.za

NOTES

¹ Systematic students of behaviour who find mutually exclusive behaviours useful can handle this with book-keeping, counting 'walking and not chewing', 'not walking and chewing' and 'walking and chewing' as three distinct behaviours.

² We could also think of preferences as representing the values of outcomes or states of affairs. Since we're talking about affordances here, I'm mostly going to consider preferences as indexed to behaviours given current state (of world and body).

³ Some sea slugs are capable of some kinds of learning, so it's important not to think that *all* of their behavior selection processes are like the one described. See Pery, Barron and Cheng (2013), for a survey of learning types among invertebrates.

⁴ For more on common currency claims in the cognitive and behavioural sciences see Spurrett (2014).

⁵ A point I don't have space to develop here is that scarcity will mean that not all affordances will be discriminated or processed in any detail, and that which ones are could itself be subject to competition and sensitive to needs and priorities. Cisek (2007) is alert to this, as is Millikan (2017).

REFERENCES

- BROOKS, R.A. (1991), 'Intelligence Without Representation', *Artificial Intelligence*, 47, pp. 139-159.
- CHEMERO, A. (2009), *Radical Embodied Cognitive Science*, Cambridge, MIT Press.
- CISEK, P. (2007), 'Cortical Mechanisms of Action Selection: The Affordance Competition Hypothesis', *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362, pp. 1585-1599.
- CLARK, A. (1997), *Being There*, Cambridge, MIT Press.
- (2016), *Surfing Uncertainty*, Oxford, Oxford University Press.
- DAVIS, W.J., MPITSOS, G.J., PINNEO, J. M., and RAM, J. L. (1977), 'Modification of the Behavioral Hierarchy of Pleurobranchaea. I. Satiation and Feeding Motivation', *Journal of Computational Physiology*, vol. 117, pp. 99-125.
- DENNETT, D.C. (1975), 'Why the Law of Effect Will Not Go Away', *Journal for the Theory of Social Behaviour*, vol. 5, pp. 169-188.
- (1991), *Consciousness Explained*, Boston, Little, Brown.
- (1995), *Darwin's Dangerous Idea*, New York, Simon & Schuster.
- (1996), *Kinds of Minds*, New York, Basic Books.
- (2017), *From Bacteria to Bach and Back*, New York, Norton.
- (2018), 'Reflections on Peter Godfrey-Smith', in Huebner, B. (ed), *The Philosophy of Daniel Dennett*, Oxford, Oxford University Press, pp. 250-253.
- DORRIS, M.C., and GLIMCHER, P. W. (2004), 'Activity in Posterior Parietal Cortex is Correlated with the Subjective Desirability of An Action', *Neuron*, vol. 44, pp. 365-378.
- GIBSON, J. J. (1979), *The Ecological Approach to Visual Perception*, Boston, Houghton Mifflin.
- GLIMCHER, P. W. (2003), 'The Neurobiology of Visual-Saccadic Decision Making', *Annual Review of Neuroscience*, vol. 26, pp. 133-179.

- (2011), *Foundations of Neuroeconomic Analysis*, Oxford: Oxford University Press.
- GODFREY-SMITH, P. (2018), 'Towers and Trees in Cognitive Evolution', in Huebner, B. (ed) *The Philosophy of Daniel Dennett*, Oxford, Oxford University Press, pp. 225-249.
- KLEIN, J.T., DEANER, R.O., and PLATT, M. L. (2008), 'Neural Correlates of Social Target Value in Macaque Parietal Cortex', *Current Biology*, vol. 18, pp. 419-424.
- LEVY, D. J. and GLIMCHER, P. W. (2012), 'The Root of All Value: A Neural Common Currency for Choice', *Current Opinion in Neurobiology*, vol. 22, pp. 1027-1038.
- MILLIKAN, R.G. (2017), *Beyond Concepts*, Oxford, Oxford University Press.
- PERRY, C. BARRON, A. and CHENG, K. (2013), 'Invertebrate Learning and Cognition: Relating Phenomena to Neural Substrate', *Wiley Interdisciplinary Reviews: Cognitive Science*, vol. 4, pp. 5610582.
- PLATT, M. L., and GLIMCHER, P. W., (1999), 'Neural Correlates of Decision Variables in Parietal Cortex', *Nature*, vol. 400, pp. 233-238.
- SHERRINGTON, C.S. (1906), *The Integrative Action of The Nervous Systems*, New Haven, Yale University Press.
- SHIZGAL, P., and CONOVER, K. (1998), 'On the neural computation of utility', *Current Directions in Psychological Science*, vol. 5, pp. 37-43.
- SPURRETT, D. (2014), 'Philosophers Should Be Interested in 'Common Currency' Claims in the Cognitive and Behavioural Sciences', *South African Journal of Philosophy*, vol. 33. pp. 211-221.