

On the Breadth and Significance of Niche Construction: A Reply to Griffiths, Okasha and Sterelny

KEVIN N. LALAND^{1,*}, JOHN ODLING-SMEE² and
MARCUS W. FELDMAN³

¹*School of Biology, St. Andrews University, UK;* ²*Institute of Biological Anthropology, Oxford University, UK;* ³*Department of Biological Sciences, Stanford University, USA; (*Author for correspondence; e-mail: knl1@st-andrews.ac.uk)*

Received 25 November 2004; accepted 26 November 2004

Key words: niche construction, evolution, ecological inheritance, ecosystem engineering, inheritance, semantic information

Introduction

We thank Paul Griffiths, Samir Okasha, and Kim Sterelny for devoting what must have been considerable time, energy and care to scrutinising our book – we are quite aware that it is not always an ‘easy read’ - and for their thoughtful and stimulating remarks. We sense that all three ‘see the big picture’ and we are greatly encouraged by their enthusiasm and positive comments. They raise many interesting issues, but the broadly supportive tone of these commentaries encourages us to believe that the niche-construction perspective has a strong theoretical foundation.

The major questions, or reservations, raised relate to (1) the breadth and inclusive nature of our conceptions of ‘niche construction’ and ‘ecological inheritance’, (2) our emphasis on ‘genetic information’ and ‘semantic information’, (3) our claim that niche construction is a ‘selective’, ‘predictive’ and ‘profitable’ process, (4) ‘positive’ and ‘negative’ niche construction, and how they associate with fitness, and (5) whether ‘The Extended Phenotype’ offers a more compelling alternative.

The breadth of ‘niche construction’ and ‘ecological inheritance’

Niche construction

All three reviewers are worried that our conception of niche construction is too broad. This concern is manifest at two levels. First, the reviewers feel that there

are qualitatively different kinds of niche construction and a more detailed taxonomy of forms is required. Second, they suggest that some phenomena that we claim to be examples of niche construction might better be called something else. That is, they suggest we have exaggerated the ubiquity and significance of niche construction, and that some phenomena we describe as ‘niche construction’ are less central, important, or consequential to the evolutionary process than what the commentators would define as ‘true’ niche construction.

We sympathise with the call for a more detailed categorization of niche construction, and recognise that we did not complete this job. However, having struggled with it over the past 15 years, we are well aware of some of the inherent difficulties involved.

Central to the niche-construction perspective lies the relativistic concept of the niche, from which a full taxonomy of niche construction would have to be derived. We believe that such a taxonomy should include at least the following classifications: (i) the different kinds of ways organisms can change the organism–environment relationship (for instance, by perturbing their environment, or relocating to a different environment), and (ii) the different kinds of organism–environment interactions (for instance, counteractive or inceptive niche construction that, respectively, negate or initiate selection). In addition, (iii) we can envisage circumstances under which it might be useful to distinguish between various types of consequences of niche construction, for instance, (a) between ecological and evolutionary consequences, (b) between consequences for the niche-constructing organisms themselves and for other organisms, (c) between consequences for environments, niches and ecosystems, (d) between short-term and long-term effects. The relative importance of these distinctions will vary from case to case.

However, when it comes to the concern that our general conception of niche construction may be too broad, we wish to stand our ground. Let us deal, one by one, with the reasons for narrowing ‘niche construction’ put forward by the reviewers. Okasha places emphasis on the distinction between activities that alter the niche of the constructor and those that affect other organisms. He would like us to restrict use of the term ‘niche construction’ to the former and gives a number of reasons why he thinks such a narrowing of the term would be a good idea. We disagree.

First, Okasha asserts “the language of construction applies more naturally to cases where organisms modify their own selective environment” ... “constructing seems to imply deliberately modifying one’s environment to suit one’s ends”. This latter claim would require a further restriction of the term to cases of niche construction that are ‘deliberate’ and ‘suit one’s ends’, that is, where there is some kind of intentionality to increase fitness.

We feel that it is important not to let the cart lead the horse here. If “the language of construction” is misleading in this regard then we would rather change our language than our argument, although we note that terms synonymous to ‘niche construction’ (e.g. ‘ecosystem engineering’) have also run into

similar criticism. Any subdivision of the process of niche construction must be justified on conceptual grounds, not because it fits better with a narrow conception of the ‘niche construction’ label.

Okasha continues “if [photosynthetic] bacteria can be said to “niche construct” it is surely in a quite different sense from that in which dam-building beavers niche-construct.” Okasha’s concern stems from the observation that “the environmental impact is a long way downstream of the organisms’ activities”. This seems to us more of a quantitative than a qualitative distinction between beavers and bacteria. At the *per capita* level, beaver’s niche construction modifies their world a great deal, to immediate effect, while niche construction by a single bacterial cell modifies the environment a small amount, with effects that are only detectible down the line, once they have accumulated. Yet researchers studying ‘ecosystem engineering’ have documented how organisms with small individual impacts can have huge ecological effects, provided the organisms occur at sufficiently high densities, or the effects accumulate over sufficient periods of time (Jones et al. 1994). For instance, to the extent that microorganisms are superabundant and share genes that cause them to express the same niche-constructing activities (e.g. photosynthesis) over long spans of time, they typically act as extremely powerful, uni-directional biological pumps for specific nutrients and gases (e.g. CO₂), with vast environmental consequences. Bog-forming *Sphagnum* mosses provide another example. Accumulated *Sphagnum* peat may persist for hundreds to thousands of years after the death of the living moss (Tansley 1949), producing a substantive and long-lasting change in selection pressures for resident populations.

In between beavers and bacteria we have earthworms. Their burrowing activities are simultaneously ‘beaver-like’, in that each worm directly benefits from its own activities, and ‘bacteria-like’, in that their impact on the soil accumulates over many generations. For instance, being structurally poorly adapted to life on land (Turner 2000), the earthworm’s burrowing and associated activities are physiologically vital, since they weaken soil matric potentials, allowing the organism to draw water into its body, thereby preventing desiccation (Turner 2000). The resultant changes in soil chemistry and nutrient cycling occur only as an accumulated effect of many generations’ activities (Lee 1985), but almost certainly feed back to modify selection pressures on descendant populations. Since instances of niche construction that are neither deliberate nor obviously beneficial to the constructor can nevertheless direct its subsequent evolution, which is our major focus, Okasha’s proposed narrowing seems to us unnecessary.

We anticipate that a single niche-constructing activity will frequently generate multiple forms of feedback over a range of temporal scales. We see no reason to focus solely on the immediate effects. This is as true for bacteria as for earthworms. For instance, gut bacterial strains can facilitate their own invasion of a new host through environmental perturbation, such as the release of bacteriocins, in spite of the presence of locally adapted resident strains (Brown et al. unpublished manuscript; Riley and Wertz 2002).

Now a single niche-constructing activity might have multiple consequences, for many organisms, and species, at various times and places. For example, the activity might generate feedback in the form of a modified natural selection pressure for the niche-constructing organisms themselves. This can be subdivided into feedback that changes the fitness of the allele(s) expressed in the niche construction and/or feedback to other genetic loci in the same organism. It can also affect selection acting on other organisms in its own population, on other populations of the same species, on other species, on descendant populations of same or different species, and so forth. So, although it may be desirable to distinguish between the consequences of niche construction that feed back to self versus those that affect “others”, in practice niche construction doesn’t work like this. This need not rule out using working definitions to indicate which kind of niche construction, and what kind of consequence, is relevant for any given case. At this level, the distinctions emphasized by Okasha and Sterelny could be valuable.

Second, Okasha suggests that it is only when organisms alter their own environment that they generate an organism–environment (O–E) match. Not so. Consider, for example, the numerous species of plants engaged in ecological successions, each of which alters its environment in a manner that enhances its successor’s match with the same environment. Ragweed constructs a soil that favours asters and broom sedge, which in turn favour blackberry, and so forth. Or take the gopher–grass–grasshopper relationship (Huntly and Inouye 1988). Gophers’ digging engineers the soil, changing vegetative composition, biomass and productivity, and favouring the plant species that gophers prefer to eat. This enhances the O–E match for gophers, but it also enhances the match for grasshoppers that exploit the same plants, and which become abundant in the vicinity of gopher mounds. There are many such examples (Jones et al. 1994, 1997; Odling-Smee et al. 2003).

Third, Okasha claims that “some of OLF’s own arguments seem to presuppose the narrower rather than the broader notion of niche-construction”, suggesting that our perturbation-relocation and inceptive-counteractive dichotomies only makes sense relative to the constructor. We think a more useful distinction here is between ‘phenotype’ and ‘extended phenotype’ (Dawkins 1982). The constructing activity (phenotype) can be described as perturbatory or relocatory, inceptive or counter-active, but the change in the environment (extended phenotype) cannot. To the extent that other organisms typically experience the change rather than the act of changing then, as Okasha says, these sub-categories of niche construction do not pertain to the modified environment of other organisms. However, neither do they relate to the modified environment of the constructor. The distinction is between constructing and construction, not between feedback to self or other.

Okasha continues “[t]heir point about the two routes to establishing an O–E match, about counteractive niche-construction shielding genetic variation from selection, about cultural niche-construction in humans interacting with genetic evolution, about the importance of ecological inheritance, about organism–

environment co-evolution, and many others, all make most sense where niche-construction is understood narrowly.” We disagree. In addition to succession, consider the case of ‘by-product mutualism’ where by-products drive co-evolutionary events, the regulatory behaviour of nest builders that shields genetic variation in inquilines, genetic responses in domesticated plants and animals to human cultural niche construction, ecological inheritance (e.g. of a beaver’s dam) by other species, and so forth. These often neglected, often indirect, multi-species interactions are likely to play critical roles in both co-evolutionary dynamics and ecosystem functioning. Okasha says “I suspect that OLF’s oscillation between the broader and narrower notion of niche-construction stems from their desire to integrate evolutionary and ecological concerns under a single theoretical structure”. He is correct that such a desire has been central to our thinking, although throughout our studies on niche construction we have been consistent in utilising the broad definition.

This brings us to a fourth line of argument, introduced by Sterelny, which distinguishes between “the mere effects of an organism’s actions on its environment, and the ability of agents to control their own environment, albeit partially”. We assume such putative ‘control’ is the product of natural selection, and hence that Sterelny is distinguishing between ‘adaptations’ and ‘effects’, *sensu* Williams (1966). We agree such a distinction is often important, and we can all point to past writings of our own in which we have been critical of others who too readily assume characters are adaptations. From the conventional perspective Sterelny is right: beaver’s dams are probably qualitatively different from the impact of micro-organisms on nutrient cycling, to the extent that the former but not the latter are adaptations.

If the only feedback to an organism from a niche-constructing activity were due to effects on selection of the genes that underpin the activity, then whether the character is an adaptation or effect is of paramount importance, since the difference between these impinges on survival value and reproductive benefits of the character. But, as all three commentators seem to accept, this is *not* the only form of feedback from niche construction. Such activity frequently also modifies selection pressures acting on other aspects of the phenotype, in the same or in descendent populations; for this second kind of feedback the distinction between adaptation and effect is irrelevant. One of the contributions of the niche-construction perspective is to focus on the symmetry between these rather than their sequential nature, which is the old way of thinking about evolution (Lewontin 1983).

One of our major points is that certain important forms of feedback in evolution are consistently neglected because the conventional perspective discourages their consideration. Feedback from “mere effects” (sometimes called “by-products” or “changes”) that are not adaptations (or “extended phenotypes”), or those stemming indirectly from other biotic, or even abiotic, components in the constructor’s niche, can be important in evolution – every bit as consequential as feedback from what are seen conventionally as adaptations. Sterelny’s use of the adjective “mere” to describe “effects” is common

within evolutionary biology, and a good illustration of the current habit of dismissing the feedback from effects as inconsequential. When earthworms change soil chemistry they modify selection pressures as surely as when they line their burrows with leaves, despite the fact that probably only the latter instance can be described as an adaptation. If we narrow the term niche construction we risk neglecting important forms of feedback stemming from effects.

We do, however, accept Sterelny's point that niche-constructing adaptations can be shaped by selection in a manner that niche-constructing effects will not. That is what gives, for instance, weaverbirds' nests the rich structure and complexity lacking in, say, detritus. Once again, however, we feel that Sterelny's focus on the magnitude of *per capita* effects, and on single versus mass effects, is misplaced. The 'proactive' gut bacteria that poison their competitors have tiny *per capita* effects that are only manifest through a mass action, which may benefit each cell no more than other cells of the same strain located in the same region, yet their bacteriocidal niche construction has been favoured by selection (Brown et al., unpublished manuscript).

Sterelny's distinction between collective and individual niche-constructing effects is also not quite as compelling as it might, at first sight, appear. Sterelny accuses us of frame shifting between individual and population effects because we lump together beavers and termites, which have big effects on their individual worlds, with soil manufacturers which exert small effects that are important only when accumulated across a population. Only the former, he suggests, have design properties. However, the natural world is not quite as simple as that. Termites nests are built by colonies not individuals. Beavers dams are built by families of beavers. Most spider's webs are individual constructions, but colonial spiders built communal webs that have lost none of their effective "design." The burrows of many mammals are built, maintained and defended by more than one individual. What we suspect Sterelny really means is not whether the particular niche construction has large *per capita* effects, or whether the construction is carried out by an individual or a population, but whether it is an adaptation.

More centrally, we take issue with Sterelny's claim that while "[m]ere effects are often important.... they do not establish the evolutionary and ecological cascades that are the consequence of life in a partially designed world." Consider, for example, the case of the Negev desert snails that graze on endolithic lichens that grow under the surface of limestone rocks (Shachak et al. 1987). The function of their herbivory is manifestly not soil weathering, which is indisputably an "effect" rather than an adaptation. Yet the snails are major agents of rock weathering and soil formation in the desert. These snails must physically disrupt and ingest the rock substrate in order to consume the lichens. They later excrete the rock material ingested as faeces, and the annual rate of biological weathering of these rocks by snails is approximately 1 metric ton per hectare per year, which is sufficient to affect the whole desert ecosystem (Shachak and Jones 1995; Shachak et al. 1987, 1995). By converting rock to

soil at this rate, the snails become major agents in soil formation, generating a cascade of ecological consequences for the distribution of desert plants and animals. In contrast, their consumption of lichen, which could legitimately be described as an adaptation, has trivial ecological consequences.

There are many such examples. “Mere effects” are every bit as likely to establish evolutionary and ecological cascades as adaptations. It is difficult to conceive of an adaptation in any species of living creature that has been as evolutionarily consequential as oxygen production by bacteria. Indeed, the effects of photosynthesizing micro-organisms would be better described as “gargantuan” than “mere”!

Ecological inheritance

Sterelny also suggests that our portrayal of ecological inheritance is too broad. He writes “[a]s I see it, Odling-Smee, Laland and Feldman’s conception of ecological inheritance amalgamates three very different intergenerational evolutionary effects”, which we summarise as (1) inherited developmental resources following relocation, (2) inherited developmental resources following perturbation, and (3) engineered developmental environments. We don’t doubt that there may be utility in these distinctions, particularly for developmental biology. We are also open to the view that some of the phenomena that we characterize as ecological inheritance may not meet strict definitions of inheritance. To be honest, we have given insufficient consideration to ‘inheritance mechanisms’, and Sterelny’s challenge should be taken up. Our goal was primarily to make the case that different forms of ecological inheritance are evolutionarily consequential, a case that Sterelny apparently accepts.

Consider three closely related species of tree frog, all of which actually exist (details in Odling-Smee et al. 2003). In the summer, to avoid desiccation of the eggs, one species deposits eggs in water, another regularly urinates on an arboreal egg cluster, and a third builds a mucus bubble nest. While these three strategies, which correspond to Sterelny’s three effects, differ in many ways, setting up alternative selection pressures on the three lineages, in one important respect they are identical. All involve parental niche-constructing acts that ensure the eggs are in a moist environment, in the process negating selection for, say, water resistant eggshells or larval aestivation. The offspring of all three species experience selection pressures modified by their parents. We see no reason to conclude that one is more important than the others. Irrespective of whether all meet narrow definitions of inheritance, for us the main point is that all of these behaviours influence the evolutionary dynamic.

It may well be the case that Sterelny’s three forms of ecological inheritance produce qualitatively different types of consequence for the evolution of the organisms concerned, although we are not entirely convinced by the details of Sterelny’s analysis. For instance, we envisage circumstances under which parental habitat choice (e.g. by ovipositing insects) narrows rather than

expands the range of developmental environments experienced by offspring. At first sight, our classification of niche construction as inceptive or counteractive better explains whether the evolutionary response is likely to be directional or stabilising. It remains to be seen whether a more detailed consideration of the inheritance mechanisms will shed light on the evolutionary dynamics, and we would certainly welcome such an analysis.

Contrary to the suggestion by the reviewers that our conception of ecological inheritance is too broad, we wish to point out that it is in fact broader than they portray it. All cases of ecological inheritance have in common the inheritance of natural selection pressures that have previously been modified by ancestral niche construction, irrespective of whether the “ancestors” were biological relatives, or “ecological relatives” (by which we mean individuals that share the same habitat). In Odling-Smee’s (1988, p. 80) original description of ecological inheritance he wrote “environmental perturbations need not be transmitted by genetic relatives only. They may also be transmitted by “neighbours”, who share, or who have previously shared, some common aspect of a mutual environment. Such neighbours have to be ecologically related but need not be genetically related”. We will leave the philosophers and developmental biologists to chew over whether this is really an inheritance, but continue to stress the evolutionary importance of such legacies. Our primary point is that evolutionary processes drive environmental changes as surely as they drive organismal changes, and that some environmental changes induced by niche-constructing organisms persist for long enough to affect the subsequent evolution of organisms.

Genetic and Semantic Information

Information ‘in genes’

We endeavoured to write *Niche Construction* in the language of population biology, since population biologists were our primary audience. This community commonly conceives of information as ‘encoded in genes’ (Maynard-Smith and Szathmary 2000), and we have followed this trend. Nonetheless, we accept Griffiths’ and Okasha’ remarks that, strictly, our talk of “information encoded in DNA that specifies the adaptations of organisms” is imprecise. We believe that DNA is an especially important information store for living organisms, but accept that it is clearly not the only information relevant to transcription (and translation) and hence to development. To the extent that we have equated information with genes, and genes with DNA, we have been guilty of over-emphasising the gene at the expense of the multitude of other factors that developmental biologists are correct to stress. The reader can take it as given that our use of the term ‘genetic information’ is a short hand for DNA plus associated cellular machinery. Our only justification, if it can be described as such, is that in our ‘battle’ to convince a sceptical community of the merits of

the niche-construction perspective we have been concerned that there may be a limit to how many ‘fronts’ it is wise for us to fight on at any one time.

Griffiths says “I found it puzzling that a book whose subject is the multiple causal channels by which a parental generation can influence their descendants would ignore the role of those very channels in the transmission of the structures which allow organisms to defeat entropy and stay alive.” He goes on: “[o]ne explanation may be that the book’s focus is so much on how organisms construct the *selective niche* of their descendants, rather than on how they construct their *developmental niche*.”

Let us try to alleviate some of Griffiths’ puzzlement. The subject of our book is not “the multiple causal channels by which a parental generation can influence their descendants”; nor is it a book about development. *Niche Construction* is concerned with the nature of the process of evolution, with the causal basis of the organism-environment match and its inherent symmetries, as well as the active role that organisms play in driving evolutionary and co-evolutionary events. Our argument would hold even if there were no such phenomena as ecological inheritance, although naturally we believe ecological inheritance strengthens our case. We agree with Griffiths that “[t]here are substantial issues that urgently need to be addressed concerning the relationship between niche construction and recent work on the ecological context of development.” Indeed, we said as much in the final pages of the book. However, the concept of niche construction is so fundamental, complex and multifaceted that it could be criticized as attempting to cover too much in biology and the human sciences. We have dedicated a substantive part of our academic lives to exploring the implications of niche construction for evolutionary biology, for population and community ecology, for ecosystem ecology, for archaeology, anthropology and psychology, and had been justifiably concerned that, in tackling so many topics, we were being overly ambitious. Niche construction is too big a topic for any single human being to master. We make no apology for prioritising evolution and ecology over other areas; the implications of niche construction are at least as important here as elsewhere. We would be flattered if developmental biologists were to take up the subject and integrate development and evolutionary ecology, as framed by niche construction, in a meaningful theory.

Semantic information

Okasha asks “Granting the notion of genetic information for the moment, why assimilate it to the information acquired through ontogenetic processes? Why use the single word “information” for both?”. He wonders whether our position is based on an adherence to the philosophical tradition of evolutionary epistemology, which stresses the universal nature of natural selection (Popper 1979; Hull 1982; Plotkin 1982, 1994; Dennett 1995). On evolutionary epistemology we three authors possess slightly different views, with one of us (JOS)

broadly supportive and the other two (KNL, MWF), while interested and sympathetic, remain sceptical. As a result, in the book we are non-committal on whether learning and other non-genetic processes of information gain are Darwinian. However, we share the belief that there are multiple information-gaining evolutionary processes, operating at different levels and with different storage mechanisms. We also agree that much of the “information” gained at distinct levels, while it may differ in many respects, is characterised by some semantic quality. Our belief stems from the observation that characters such as the capacity to learn, or the vertebrate immune system, are adaptations, and that particular cases of learned behaviour or immunological response may well be adaptive. It follows that the information accrued through these systems relates functionally to specific features of organisms and specific aspects of their environments, and therefore that it conveys *meaning*. We cannot rule out the possibility that there may be a qualitative difference between the kind of semantic information encoded through population genetics and through learning, but from our perspective this doesn’t matter. If these are different, and the differences are important, it would be interesting to know how and why.

A central plank of our argument is the postulate that niche construction is a universal feature of living systems and that this postulate follows from a consideration of the fundamentals of thermodynamics. This claim is challenged by Griffiths, who writes: “(i) the argument from accepting this thermodynamic characterisation of organisms to the conclusion that their activities must be guided by information is much less straightforward than the authors suppose, (ii) ... semantic information is, in any case, entirely the wrong sort of information to make the necessary connections to thermodynamics, and (iii) ... locating semantic information primarily in DNA sequence sacrifices many of the potential insights of the niche construction framework.” We have already discussed point (iii), where we are largely in agreement with Griffiths, and we point out that two entire chapters of our book were dedicated to consideration of niche construction stemming from non-genetic information. Here we focus on Griffiths’ first two points.

In Chapter 4 of our book we note that organisms are far-from-equilibrium systems relative to their physical surroundings and ask how living creatures maintain this status. Maxwell’s Demon was introduced merely as an analogical tool to address this question. Long consideration by physicists and philosophers has uncovered a number of putative properties that Maxwell’s demon would have to possess in order to perpetuate a far-from-equilibrium system. Our goal is not to prove that these putative properties of Maxwell’s demon are correct, or to consider whether they apply universally to all conceivable demons. Rather, we ask, assuming that these properties are correct, what does it tell us about niche construction? Naturally, if the putative properties are not general, then our inferences will not hold. We have used Maxwell’s Demon (in preference to, say, automata theory) merely as a guide to what we are really after, namely the establishment of the general qualitative characteristics of

niche construction. As we are not trying to advance physics, it does not matter if we have made a mistake about Maxwell's Demon (we don't believe we have) so long as we are correct about the general properties of niche construction.

Griffiths identifies six properties of a 'fool's demon': Such 'demons', "and thus organisms, must (1) have a goal; (2) discriminate states of their environment; (3) anticipate those states; and in order to do this must (4) be *instructed* by *knowledge*' (p. 173, their emphases), which in turn implies possessing a memory. They must (5) be designed in a way that allows them to accomplish these things and (6) be supplied with an external source of energy." Griffiths accepts (i), (v) and (vi) as uncontroversial, but disputes the others.

Let us focus on property 4, be *instructed by knowledge*, which appears most central to Griffiths' concerns, and which, as we specify in our book, "is the demon's most controversial property". (Properties 2 and 3 are discussed below). In fact, recent analyses of this problem specify that it must be 'knowledge' and not Shannon-like configurational information that instructs the demon (e.g. Jaynes 1996). Only this kind of information carries 'meaning', 'memory', and 'anticipates', properties that the demon must have to function effectively. For physicists the problem is where does this knowledge come from? For biologists there is no such problem; knowledge stems from natural selection.

For us, as for most biologists, any kind of adaptive trait that is at least in part a product of the prior natural selection of genes is both evidence for, and expression of, semantic information. Such organismal adaptation is, of course, not exclusively dependent on the inheritance of "remembered" semantic information in previously selected genes. However, it must start there.

Why, then, does Griffiths assert that semantic information is "entirely the wrong sort of information to make the necessary connections to thermodynamics"? For Griffiths "the problem is that...the relevant concept of information is quite different from that which figures in the mathematical theory of communication and it is only the later concept that has a formal connection to the measurement of entropy and thus to thermodynamics.... The 'knowledge' or 'meaningful information' that flows from parent to offspring is simply not the same thing as the quantity of Shannon-information that a descendant genome carries about an ancestral genome". We agree; indeed throughout our book we drew the same distinction, at some length.

Where we differ from Griffiths is in his conclusion that "it seems reasonably clear to me that it is Shannon information that must be the focus of any attempt to use natural selection as the external source of order that makes organisms Fool's demons.... This because it is Shannon information that will reflect the imposition of order on one physical system by a causal connection to another physical system and... it is by this means that natural selection is able to construct dissipative systems."

Two steps are needed to respond to this assertion. First, we maintain that semantic information *is* the right kind of information to talk about within biology. On this we are in broad agreement with Maynard Smith's (2000)

seminal discussion of the concept and, we believe, with the general tenor of current biological thought. Second, it follows that, if we are right to draw parallels between the putative properties of Maxwell's demon and of niche-constructing organisms (and we still think we are), then it must be possible to make some kind of connection between semantic information and thermodynamics. But how?

The trouble is that, as far as we are aware, nowhere in science have all the issues relating to either "information" or "thermodynamics" been settled. Indeed, physicists are still arguing about what information and entropy are (Smolin 2001; Bekenstein 2003). Little wonder that mere biologists struggle with these issues.

Consider the following example (Laughlin, personal communication), which illustrates some of the relevant difficulties. If individuals want a new passport they have to fill in a questionnaire asking them whether they are male or female. When they tick the box labelled "sex" their answering tick supplies a single binary digit of information. But that can't be all there is to it because, in this case, one BIT of information is clearly referring to a huge amount of semantic information too. Whether an individual is male or female certainly is meaningful in our niche. To our knowledge no one has yet managed to spell out exactly how it is possible to collapse so much semantic information into one binary digit. Like others, we think, this "trick" is achieved by vast amounts of prior (naturally) selective deaths (Kimura 1961; Odum 1988), plus vast amounts of prior developmental and/or socio-cultural processing, affecting and informing and "preparing" the lives of individual organisms. It is only this preparedness that ultimately makes the "tick" so meaningful to the individual who reads the passport questionnaire.

We are happy to plead guilty to ignoring this issue completely in our book, where we merely distinguished between configurational (Shannon) information and semantic information. We did not attempt to relate the one to the other, in any formal manner, because we didn't know how to do it. As far as we know, nor does anyone else.

But why do we need a formal connection here? We are not trying to derive a mathematics of niche construction from a mathematics of thermodynamics – our goal is more modest. It may even be possible for prior knowledge to be a prerequisite for any system that will counteract entropic processes without there being any formal connection between 'knowledge' and 'thermodynamics'. Our argument that there is likely to be a connection between semantic information and thermodynamics operates at a more intuitive level.

First, while adaptive evolution is possible because natural selection is a process that accrues semantic information, as a direct consequence of this, Shannon-type information must also accumulate, if for no other reason than because all knowledge requires a physical store. There is no direct correspondence between knowledge-gain and Shannon-type information gain, and in many instances increases in one will be associated with decreases in the other. However, there must often be a loose correlation, since knowledge

requires configurational information, and the amount of configurational information that can be stored imposes minimal constraints on the amount of knowledge that can be stored. At its upper limit, an organism with a large genome is potentially more complex than an organism with a small genome. The same applies to brains and the complexity of behaviour (Attwell and Laughlin 2001). The observation that salamanders have huge genomes relative to humans, or that junk DNA exists, is testament to the fact that the correlation between the amount of semantic and configurational information is weak – it does not follow that the correlation is absent.

Second, organisms feed on molecules rich in free energy, in the process generating outputs largely in the form of molecules that are poorer in free energy. They use this energy to do work. Such work is necessary to allow organisms to produce and maintain order, be it inside their bodies or in their external environments (Turner 2000). In order to survive, organisms must act on their environments and by doing so they must change them. Work in the thermodynamic sense is only done when energy is made to flow. Hence, work is the processes that organisms engage in that allows them to exchange energy with their environments, to channel energy through their bodies and to create orderliness in the process. Such work is a necessary condition for organisms to survive and maintain their far-from-equilibrium status against entropic forces. Niche construction is connected to thermodynamics by the fact that it *is* work.

Niche construction is a ‘selective’, ‘predictive’ and ‘profitable’ process

We sketched the reasoning leading us to the claims that niche construction is selective, predictive and profitable in chapter 4 of our book, but only briefly. This may have led Sterelny and Okasha to believe we were making stronger claims than we intended. We see little genuine disagreement here.

By “selective process” we mean only that the niche-constructing acts of organisms depend on streams of non-random developmental, physiological, functional and behavioural “decisions” (using the word “decision” as in “statistical decision theory”) throughout the lives of individual organisms. These decisions are non-random because they are partly controlled by semantic information, derived from a multitude of sources including naturally selected genes, epigenetic processes, immune systems, central nervous systems and, as in humans, cultural processes. Such a claim may be regarded as sufficiently obvious to be trite, but we did not intend that it be perceived as an important insight, it is not a major part of our argument, and we make nothing of it. It is merely included for completeness.

Our argument that niche construction is predictive is also intuitive, and once again the reviewers appear to be reading more into this than we intended, which is merely that inherited semantic information has the crucial job of “preparing” organisms in advance for their niches, and functions in this manner to the extent that niches retain some degree of autocorrelation, such

that characters selected in ancestors are of utility to descendants. We mean that for living organisms to survive they must *a priori* possess relevant semantic information, like Kant's *a prioris*, except that here we are talking about evolutionary rather than God-given knowledge. Semantic information must be associated with adaptive character states, at least more often than not. Such characters allow organisms to discriminate niche states *a priori*, or at least allow them the plasticity to develop powers of discrimination ontogenetically during their lives. Without such discrimination, living creatures would behave, grow, even metabolise randomly – that is, their niche construction would be non-selective. However, it must be based on semantic information if the properties of living organisms are to be functional and adaptive.

As Sterelny aptly puts it, “[t]he design of a filter-feeder’s filtering apparatus is an inductive bet that future environments will resemble past ones”. In describing niche constructive as predictive we mean only that it is informed *a priori* by predictive information, and not that the organism is inevitably consciously, deliberately or actively predicting its world. We agree with Sterelny that, depending on how it is defined, there may be “no search plan at the level of individual agency”. Sterelny seems to find this position unsatisfactory, asserting that “[a] notion of individual agency that counts filter-feeders ... as actively searching their environment or predicting its future is too weak”. Too weak for what? It is not ‘too weak’ for the legions of social scientists that regard human and animal behaviour as uninformed by evolved *a prioris*. Nor is it ‘too weak’ for the many biologists who fail to consider why such *a prioris* might exist. Moreover, we do not rule out the possibility that niche construction is, in restricted circumstances, ‘predictive at the level of individual agency’, and are willing to assert that, where it is so, it is solely because that niche construction is based on semantic information, often non-genetic in character.

Turning to the issue of the profitability of niche construction, Sterelny writes: “I do not think it follows that because actions must in general be profitable, organisms must typically be shaping their environment in ways that improve the organism–environment fit.... It is true that actions must indeed normally be profitable. But they are profitable not because of their effects on the agent’s environment, but because of their effects on the agent itself”. We agree. We have been quite clear (1) that niche construction can have both positive and negative effects on an organisms fitness, and (2) that it can both enhance and degrade the fit between organism and environment, but claims (1) and (2) are quite different. While we anticipate that most niche construction is positive (increasing the short-term fitness of the constructor) we do not claim that most niche construction improves an organism’s fit with its environment. On the contrary, where organisms degrade their environment it is usually because it is profitable to do so in the short-term. Sterelny seems to read us as suggesting that positive and negative niche construction directly equate to an improvement or degradation in the O–E relationship, respectively. This is not correct, although for mathematical convenience we have incorporated this

assumption in our formal models (see below), which may be the source of Sterelny's concern. These are two independent dimensions.

Sterelny continues "Odling-Smee, Laland and Feldman's own thermodynamic perspective implies that actions do not typically create an ordered environment. On average, life pumps entropy into the environment making it, if anything, worse." Here we do take issue. While at a global level Sterelny has to be right, we maintain, however, that niche construction frequently creates an ordered *local* environment. That is, niche construction creates order in the constructor's niche, and can even impose "design" on artefacts, while generating disorder elsewhere, or perhaps creates order in one part of its environment and disorder in another part.

Sterelny claims "[t]hermodynamic considerations ... indicate that these [niche-constructing] impacts are typically *mere effects*: they make the environment more disordered and less friendly, rather than consist in an extension of the agent's control over aspects of its world... Thus the ordinary ecological life of an agent has effects quite unlike those of termite mounds and the other paradigm cases of niche construction with which we began.... Agents will not typically be able to impose their own stamp on their biological world, any more than a pub after an all-in brawl reflects any agent's control." There are several respects in which we disagree with this statement.

First, as mentioned above, there may well be a useful qualitative distinction between niche-constructing adaptations and effects, but the latter are every bit as consequential as the former. We strongly dispute any suggestion that only the former category matters in evolution. Second, we contest any implication that the category of niche-constructing adaptations is small, such that agents rarely exert control over aspects of their world. Consider, for instance, the relocalational sub-category of niche construction, which we define as occurring "when organisms actively move in space, as well as choose or bias the direction, and the distance in space through which they travel, and the time when they travel, thereby modifying natural selection pressures." All motile organisms exert some control on their selective environments by virtue of the fact that they can relocate to places that suit them better. To pick up on Sterelny's analogy, if we don't like the brawl we can leave the pub. Take Sterelny's example of a filter-feeder, which he champions as an organism that lacks control in constructing its niche. There is no such absence of control for all the whales, fish, insects and other motile filter feeders that actively select where they feed (Odling-Smee et al. 2003). It is not even absent in some of the simplest animal filter feeders, the protozoa, which generate a water current with their cilia, sending food their own way (Larousse Encyclopaedia of Animal Life 1973). A closer look at the diversity of niche construction in the natural world reveals much more 'control' than might be assumed. Third, niche-constructing effects can produce order as well as disorder. For instance, as we have already mentioned, many intra- and inter-specific mutualisms have evolved because neighbours first exploit, then invest in, the by-products produced by organisms (Connor 1995).

Positive and negative niche construction

Okasha detects some ambiguity in our use of the terms ‘positive’ and ‘negative’ niche construction. We agree there is some slippage in how we use these terms, particularly in relation to our mathematical models. We define positive and negative niche construction as niche-constructing acts that, on average increase or decrease the fitness of the constructor. We change how we formalise this in going from model 1 (where positive and negative are represented by the sign of ε) to model 2 (where positive and negative are represented by λ_2 and γ). In both cases, for simplicity, we assume the resource is a valuable commodity so positive values of coefficients that result in an increment or decrement can be characterised as positive or negative niche construction respectively. Model 2 is more realistic than model 1 in several respects, including in how positive and negative niche construction are captured. In all cases in our models we are referring to relative fitness which is natural since we make standard population genetic assumptions, such as infinite population size.

However, Okasha has a point here, for two reasons. First, our use of ε in model 1 depicts situations in which positive and negative values of ε represent resource accumulation and depletion. Given the lack of spatial structure in our models, any increment in R benefits the entire population, and this formulation in a population of finite size would affect absolute fitness. Second, in the glossary we say “In the long run we expect some niche-constructing activities of organisms to become negative for their populations.” While we do expect this, for consistency we should perhaps have written “for the constructor”, which would be consistent with our focus on relative fitness.

Niche construction or extended phenotype?

Sterelny points out that Dawkins’ extended phenotype analysis raises the option of treating animal artefacts as part of the organism, not part of its world, which allows the evolution of artefacts to be treated in the same manner as that of other complex adaptations. “[I]f the evolutionary dynamics of extended phenotype adaptations are just like those of other complex adaptations, then the evolution of nests and burrows is not an instance of agent’s changing their environments, it is simply an instance of adaptation to the environment.”

That it is possible to describe the world in this way is well-established through Dawkins’ writings. We are fully aware of this option and see some merit in it. However, we chose the niche-construction perspective over the extended phenotype account for three reasons.

The first is that we believe the niche-construction perspective to be a more accurate representation of causality in evolution than more conventional

perspectives, such as Dawkins'. Dawkins' solution is both simple and elegant, but there is a price to pay for it, in that it neglects certain forms of feedback and narrows the evolutionary focus. There are two major problems with Dawkins' position, on which one of us has elaborated elsewhere (Laland 2004). First, as our models show, the evolutionary dynamics of extended-phenotype adaptations are manifestly *not* just like those of other complex adaptations. We have documented many unusual evolutionary dynamics (timelags, inertia, momentum, etc.) that result primarily because extended-phenotype adaptations and other niche-constructed effects have carry-over effects from one generation to the next (Laland et al. 1996, 1999, 2001; Odling-Smee et al. 2003). The second, and more fundamental, is that the conventional view promotes the unfortunate tendency to assume that only characters that are adaptations are evolutionarily important. For instance, Dawkins' approach neglects niche construction resulting from by-products and other non-adaptations, which can equally be consequential. Also, once we recognise that there is a second route by which phenotypes play a role in evolution, and a second form of feedback from niche-constructing effects, it opens the door for a multitude of developmental processes, acquired characters, social learning and culture to be instrumental in the evolutionary process, through their influence on niche construction. For example, it grants phenotypes a limited capacity to co-direct the genetic evolution of their populations by recruiting ontogenetic processes to modify natural selection. That raises philosophical issues that are more often associated with "Lamarckism". However, niche construction is not Lamarckian, It is Darwinian. It only modifies orthodox Darwinian selection.

The second reason is our desire to focus on the symmetry between organism and environment that is the natural qualitative consequence of Lewontin's (1983) original pair of differential equations. The extended-phenotype approach retains the notion of organismal adaptation to environment rather than the coevolutionary dynamics of both.

The third reason for our preferring the niche-construction approach, perhaps more significant to empirical researchers than to philosophers, is that it brings with it a fresh perspective, with a suite of novel hypotheses and methods, some of which we have begun to catalogue (Odling-Smee et al. 2003). One of our principle objectives, in writing *Niche Construction*, was to take Lewontin's original idea and to develop it into a viable empirical programme of research. We hope and believe that the approach is likely to stimulate researchers to contribute to this program, and ultimately feel that it is here that the perspective will stand or fall. We were surprised that none of the reviewers explicitly questioned our claim that niche construction is a second major adaptive process in evolution, and are tempted to read that as suggesting that the niche-construction perspective is a tenable one. If our supposition is correct, it would seem that the debate over niche construction is moving on to new phase, related more to utility than descriptive accuracy.

Acknowledgements

We are grateful to the NIH, The Royal Society and The Leverhulme Trust for financial support and J. Scott Turner for comments on an earlier draft. We would also like to thank Kim Sterelny for commissioning this special edition. We are proud and grateful that our book has been singled out in this manner.

References

- Attwell D. and Laughlin S.B. 2001. An energy budget for signalling in the grey matter of the brain. *J. Cerebr. Blood F. Metabol.* 21: 1133–1145.
- Bekenstein J.D. 2003. Information in the holographic universe. *Sci. Am.* (August): 59–65.
- Brown S.P., Le Chat L. and Taddei F. (unpublished manuscript). Autocatalytic environmental perturbation allows rare invaders to replace adapted residents.
- Connor R.C. 1995. The benefits of mutualism: a conceptual framework. *Biol. Rev.* 70: 427–457.
- Dawkins R. 1982. *The Extended Phenotype*. Freeman, Oxford.
- Dennett D. 1995. *Darwin's Dangerous Idea. Evolution and the Meanings of Life*. Penguin books.
- Hull D.L. 1982. The naked meme. In: Plotkin H.C. (ed.), *Learning, Development, and Culture*. Wiley, Chichester pp. 273–327.
- Huntly N. and Inouye R. 1988. Pocket gophers in ecosystems: patterns and mechanisms. *Bioscience* 38: 786–793.
- Jaynes E.T. 1996. The Gibbs paradox. In: Smith C.R., Erickson G.J. and Neudorfer P.O. (eds), *Maximum Entropy and Bayesian Methods*. Kluwer Academic, Dordrecht, Holland.
- Jones C.G., Lawton J.H. and Shachak M. 1994. Organisms as ecosystem engineers. *Oikos*. 69: 373–386.
- Jones C.G., Lawton J.H. and Shachak M. 1997. Positive and negative effects of organisms as physical ecosystem engineers. *Ecology* 78: 1946–1957.
- Kimura M. 1961. Natural selection as a process of accumulating genetic information in adaptive evolution. *Gen. Res.* 2: 127–140.
- Laland K.N. 2004. Extending the extended phenotype. *Biol. Phil.*
- Laland K.N., Odling-Smee F.J. and Feldman M.W. 1996. On the evolutionary consequences of niche construction. *J. Evol. Biol.* 9: 293–316.
- Laland K.N., Odling-Smee F.J. and Feldman M.W. 1999. Evolutionary consequences of niche construction and their implications for ecology. *Proc. Natl. Acad. Sci. USA* 96: 10242–10247.
- Laland K.N., Odling-Smee F.J. and Feldman M.W. 2001. Cultural niche construction and human evolution. *J. Evol. Biol.* 14: 22–33.
- Larousse Encyclopaedia of Animal Life 1973. Paul Hamlyn. London.
- Lee K.E. 1985. *Earthworms: Their Ecology and Relation with Soil and Land Use*. Academic Press, London.
- Lewontin R.C. 1983. Gene, organism and environment. In: Bendall D.S. (ed.), *Evolution from Molecules to Men*. CUP, Cambridge, UK.
- Maynard Smith J. and Szathmari E. 2000. *The Origins of Life. From the Birth of Life to the Origins of Language*. Oxford University Press.
- Odling-Smee F.J. 1988. Niche constructing phenotypes. In: Plotkin H.C. (ed.), *The Role of Behavior in Evolution*. MIT Press, Cambridge pp. 73–132.
- Odling-Smee F.J., Laland K.N. and Feldman M.W. 2003. *Niche Construction. The Neglected Process in Evolution*. Monographs in Population Biology 37, Princeton University Press.
- Odum H.T. 1988. Self-organization, transformity and information. *Science* 242: 1132–1139.
- Plotkin H.C. (ed.) 1982. *Learning, Development, and Culture: Essays in Evolutionary Epistemology*. Wiley, Chichester.
- Plotkin H. 1994. *Darwin Machines and the Nature of Knowledge*. Penguin Books, London.

- Popper K.R. 1979. *Objective Knowledge: An Evolutionary Approach*. Clarendon Press, Oxford.
- Riley and Wertz 2002. (cited in Brown et al.) .
- Shachak M., Jones C.G. and Granot Y. 1987. Herbivory in rocks and weathering of a desert. *Science* 236: 1098–1099.
- Shachak M. and Jones C.G. 1995. Ecological flow chains and ecological systems: concepts for linking species and ecosystem perspectives. In: Jones C.G. and Lawton J.H. (eds), *Linking Species and Ecosystems*. Chapman and Hall, New York.
- Shachak M., Jones C.G. and Brand S. 1995. The role of animals in an arid ecosystem: snails and isopods as controllers of soil formation, erosion and desalinization. *Adv. Geo Ecol.* 28: 37–50.
- Smolin L. 2001. *Three Roads to Quantum Gravity*. Basic books.
- Tansley A.G. 1949. *Britain's Green Mantle*. Allen and Unwin, London.
- Turner J.S. 2000. *The Extended Organism: The Physiology of Animal-built Structures*. Harvard University Press, Cambridge, MA.
- Williams G.C. 1966. *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. Princeton University Press.