Biology and Philosophy (2005) 20:1–10
DOI 10.1007/s10539-005-0431-3

Book review

On Niche Construction and Extended Evolutionary Theory

SAMIR OKASHA
Department of Philosophy
University of Bristol
Bristol BS8 1TB
UK
E-mail: samir.okasha@bristol.ac.uk

A review of F.J. Odling-Smee, Kevin N. Laland and Marcus W. Feldman, *Comments on Niche Construction: The Neglected Process in Evolution*, 2003, Princeton University Press, 468 pp., ISBN 0-69104437-6, $39.50.

That the activities of organisms lead to environmental modifications, on both on a local and global scale, can hardly be doubted; so much is obvious to even the most casual observer of the natural world. For this very reason, the central thesis of Odling-Smee, Laland and Feldman's book – that standard evolutionary theory needs to be replaced by an 'extended' evolutionary theory which recognises the universality and significance of niche-construction for both evolution and ecology – is a startling one. For if everyone agrees that niche-construction, i.e. organism-induced modification of the environment, occurs, why does orthodox evolutionary theory fail to take account of it? Part of the answer, as Odling-Smee, Laland and Feldman (hereafter **OLF**) recognise, is that scientific theories typically aspire to a high level of generality, which is often achievable only by deliberately ignoring many of the messy details of reality. Treating the external environment as a given – an 'exogenous factor' as economists would say – helps to keep evolutionary theory relatively simple, allowing evolutionists to focus on the central task of explaining how organisms come to have the adaptations they do. In reality, the environment is not independent of organisms' activities, as everyone knows, but treating it as if it were is justified in the interest of theoretical simplicity.

The great merit of **OLF**'s book, to my mind, is their demonstration that this 'pragmatic' justification for ignoring niche-construction is not justification enough, even if it was in the early days of neo-Darwinism. Including niche-construction *does* make evolutionary theory more complicated, as the authors acknowledge, but the additional complexity is amply compensated for by the richness of the resulting theoretical structure. As the authors' population-genetic models in chapter 3 show, niche-construction and the 'ecological inheritance' that it generates can lead to novel evolutionary dynamics, quite different from those that would otherwise exist; and as they stress in chapter 6, cultural niche-construction in our own species can dramatically alter the rate and direction of human evolution, both through its own impact and via its impact on the subsequent selection of genetic variation. Taken in conjunction with the detailed evidence for niche-construction compiled in chapter 2, and the implications for ecology discussed in chapter 5, the result is a highly convincing case for 'extended evolutionary theory'. This reviewer, at any rate, was convinced. In what follows, I raise a number of queries about the conceptual, rather than the empirical or theoretical, aspects of **OLF**'s argument.

OLF define niche-construction in a highly inclusive way, to include literally *any* organism-induced modifications of the environment. Their official definition says that niche-construction is "the process whereby organisms, through their metabolism, their activities, and their choices, modify their own and/or each other's niches", where a niche is defined as the "sum of all the natural selection pressures to which the population is exposed" (p. 419). Thus activities as different as beavers building dams in their local stream, and micro-organisms causing oxygen to accrue in the atmosphere over millennia, both count as niche-construction for **OLF** – both involve organisms modifying the environment in a way that influences subsequent selection pressures. A definition this broad has its advantages, not least that it makes it impossible to deny that niche-construction actually occurs, but at times I felt it was *too* broad, including under one umbrella phenomena that belong apart. In particular, the distinction between organisms' activities which alter the niche of those self-same organisms, and activities which alter the niche of *other* organisms, seems highly pertinent from an evolutionary point of view; but activities of both types fall under the rubric of niche-construction as defined by **OLF**.

This point is important for a number of reasons. Firstly, the language of 'construction' applies much more naturally to cases where organisms modify their own selective environment, such as dam-building by beavers, web-spinning by spiders, agriculture by humans etc. In cases where the activities of one population or species modify the selective environment of another population or species, it makes less sense to think of this as niche-construction, in my view. For 'constructing' seems to imply deliberately modifying one's own environment to suit one's own ends, rather than simply having some causal impact or other on the environment. This is especially so where the environmental impact is a long way downstream of the organisms' activities – as in the example of photosynthetic bacteria affecting the atmosphere's oxygen content. If these bacteria can be said to "niche construct", it is surely in a quite different sense from that in which dam-building beavers niche construct. It might be more useful to understand niche-construction less catholically, so that not every organism-induced modification of the environment automatically qualifies as a case of niche-construction.

Secondly, **OLF** motivate their overall argument by pointing out, quite correctly, that natural selection is not the only route to establishing an O-E match, i.e. a 'fit' between organism and environment – as Lewontin first noted. In natural selection, O changes over phylogenetic time to match E, in the way described by classical Darwinian theory. But as **OLF** stress, another possibility is that organisms may alter E over ontogenetic time to suit their pre-existing phenotypic characters (O), thus establishing the O-E match by a different route. As they say, earthworms change the soil to meet the demands of their freshwater physiology, rather than their physiology changing gradually by natural selection (p. 375). That the O-E match can be established by two routes, rather than just one, is an excellent argument for according niche-construction a role alongside natural selection in evolutionary theory. But this argument applies only where organisms alter their *own* environment, not that of others. If the activities of one species modify the selective environment of another species, it is hard to see how this will contribute to an O-E match, except possibly by chance. Far more likely, it will generate an O-E *mis*match – presuming the second species was well-adapted to start with. In so far as the motivation for introducing niche-construction derives from Lewontin's point about natural selection not being the only way to generate an O-E match, the notion of niche-construction should surely be restricted to organismic activities which alter the environment of those organisms themselves. Again, **OLF**'s concept of niche-construction seems too broad.

Thirdly, some of **OLF**'s own arguments seem to presuppose the narrower rather than the broader notion of niche-construction. For example, in chapter 2 they distinguish between 'perturbative' and 'relocational' niche-construction. Perturbation occurs when organisms "actively change one or more factors in their environments at specified locations and times by physically changing them", while relocation occurs when organisms "actively move in space", thereby "exposing themselves to different environmental factors" (p. 44). It is hard to see how this perturbation/relocation distinction can apply to cases where organisms modify the selective environment of other species; the distinction seems to be precisely between two ways in which organisms can modify their *own* selective environment. Similar remarks apply to the distinction between inceptive and counteractive niche construction – as defined by **OLF**, it is a distinction between modifications that counteract an environmental degradation and ones that don't, where in both cases the environment that gets modified is that of the niche-constructing organisms themselves (pp. 45–46). Therefore, cases of niche-construction where organisms modify the selective environment of others cannot be categorised in terms of the perturbation/relocation and inceptive/counteractive distinctions. So at least sometimes, **OLF** operate with the narrower rather than the broader notion of niche-construction, their official definition notwithstanding.

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. From an ecological point of view, the most important type of organism-induced environmental modifications are surely those that affect the selective environments of *other* species – for it is these that will affect ecosystem dynamics, competitive exclusions, mutualism, predator-prey interactions, ecosystem stability, and the like. From an evolutionary point of view these environmental modifications are also highly important, of course – as the builders of co-evolutionary models have long recognised. But so too are the modifications that organisms make to their own selective environments – what I have been calling niche-construction in the narrow sense – and when **OLF** are discussing evolution rather than ecology, their focus is often on niche-construction of this sort. Moreover, many of their most interesting arguments apply only or primarily to narrow-sense niche construction. Their 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. Restricting the notion of niche-construction to the narrow reading would result in a concept of much less *ecological* significance than **OLF**'s concept, but from an *evolutionary* point of view it might make more sense, given the importance of the distinction between altering one's own selective environment and altering that of others. Perhaps the moral is that the conceptual demands of evolution and ecology are hard to satisfy simultaneously.

Turning to a different matter, in chapter 2 **OLF** draw an important distinction between 'positive' and 'negative' niche-construction, a distinction then put to use in the formal population-genetic models of chapter 3. But there seems to be a slight ambiguity concerning the precise meaning of the qualifiers 'positive' and 'negative'. In chapter 2, **OLF** define positive niche-construction as "activities that change environmental factors into states that on average increase the fitness of the niche-constructing organism" while negative niche-construction refers to “niche-constructing activities that change environments in such a way as to reduce fitness” (p. 47) However, there are two possible ways this can be taken. Do **OLF** mean that positive niche-construction increases the *absolute* fitness of a niche-constructing organism over what it would otherwise have been? Or do they mean that positive niche-construction increases the *relative* fitness of a niche-constructing organism, relative to organisms that do *not* engage in the niche-constructing activity? It is easy to see that these are non-equivalent conditions: an activity could increase the absolute fitness of the niche-constructor but simultaneously increase the absolute fitness of non-niche-constructors by even more, in which case it would satisfy the first condition but not by the second; the converse is also possible. I call the two readings the ‘absolute fitness’ and the ‘relative fitness’ criteria for positive niche-construction respectively.

OLF do not explicitly discuss the absolute/relative fitness ambiguity. But in chapter 2 they appear to employ the relative fitness criterion, for they explicitly say that virtually all niche-constructing activities must be positive in the short-term, or else they will not evolve by natural selection in the first place (p. 48) (This is compatible with the cumulative effect of a niche-constructing activity over many generations being deleterious for the population as a whole, as **OLF** note; this is simply an instance of the well-known point that what selection favours in the short-term may have a negative impact on the species in the long-term.) So **OLF** take ‘positive’ niche-constructing activities to be ones which are favoured by natural selection, at least in the short-term. Since a niche-constructing activity will only be favoured by natural selection if it confers a *relative* fitness advantage on organisms that engage in it, vis-à-vis ones that don’t, it seems that **OLF** are adopting the relative fitness reading of ‘positive’. Whether an activity increases the *absolute* fitness of the organism, or the mean absolute fitness of all organisms in the population, is irrelevant to whether it will evolve; that depends on relative fitness. In short, if ‘positive’ niche-construction means activities that are selectively advantageous in the short-term for the niche-constructing organism, as **OLF** say, then ‘positive’ must be defined by the relative not the absolute criterion.

When we examine the formal models in chapter 3, however, it appears that ‘positive’ is being defined by the absolute criterion. **OLF** present two simple models of niche-construction based on two-locus population genetic theory. In both models the **E** locus, with alleles *E* and *e*, affects the changes that organisms make to a resource **R** in their environment, while the **A** locus, with alleles *A* and *a*, has an effect on organismic viability that depends on the amount of **R**. Table 3.1 on page 137 gives expressions for the fitnesses of each of the six two-locus genotypes for Model 1, the simpler of the two models. **OLF** then say that a 'positive' niche-constructing activity is one for which $\varepsilon > 0$, while a 'negative' activity is one for which $\varepsilon < 0$ (p. 138, p. 142). However, the condition $\varepsilon > 0$ does *not* mean that the *E* allele, and hence the niche constructing activity that it causes, will spread by natural selection. The *E* allele will spread by selection if *EE* genotypes, which engage in the most niche construction, are fitter than the *Ee* and *ee* genotypes, which engage in less; the condition for this is $\alpha_1 > 1 > \alpha_2$. (This is what **OLF** call 'external selection at the **E** locus favouring the *E* allele'.) The condition $\varepsilon > 0$ means that the niche-constructing activity raises the absolute fitness of every organism, *irrespective of genotype*, over what it would otherwise have been – as an inspection of the two-locus fitnesses in Table 3.1 shows. But what determines whether the niche-constructing activity will spread is whether the *relative* fitness of niche-constructors is greater than that of non-niche-constructors, and the condition for that is $\alpha_1 > 1 > \alpha_2$, not $\varepsilon > 0$.

In effect, this means that **OLF** have switched to an absolute fitness definition of positive niche-construction in the formal models of chapter 3. There is nothing inherently wrong with defining 'positive' by an absolute fitness criterion, of course, so long as we bear in mind that so defined, it does not follow that positive niche-constructing activities will be favoured by natural selection. Where 'positive' and 'negative' are defined absolutely, by $\varepsilon > 0$ and $\varepsilon < 0$ respectively, there is no reason to think that natural selection will favour niche-constructing activities that are positive, in either the short or long-term. But in chapter 2, **OLF** argued that there *is* reason to suppose all niche-constructing activities will be positive in the short-term, since "few organisms are likely to niche-construct in ways that reduce their immediate fitness" (p. 48); this argument is only correct if 'positive' is defined by the relative criterion, rather than the absolute one. In short, there is a tension between **OLF**'s verbal definition of 'positive' in chapter 2 and their algebraic definition of 'positive' in chapter 3. If the argument of chapter 2, that selection will favour positive niche-constructing activities in the short-term, is to be sustained, then 'positive' in chapter 3 should be defined by $\alpha_1 > 1 > \alpha_2$, not by $\varepsilon > 0$. Defining 'positive' by $\varepsilon > 0$ is all very well, but consistency then requires that **OLF** abandon their earlier argument that short-term natural selection will favour niche-constructing activities that are positive.

The ambiguity between the relative and absolute definitions of 'positive' niche-construction is closely analogous to the well-known ambiguity between 'strong' and 'weak' altruism, discussed at length by D.S. Wilson (1980); see Kerr and Godfrey-Smith (2002) for a recent discussion. In Wilson's terminology, 'strong' altruism refers to actions that reduce an organism's absolute fitness, while 'weak' altruism refers to actions that raise an organism's absolute fitness but simultaneously raise the fitness of others by even more – with the result that the organism's *relative* fitness vis-à-vis non-altruists is reduced. Even though weakly altruistic actions boost the absolute fitness of their performers, weak altruism cannot spread by selection in an unstructured panmictic population, for the relative fitness of weak altruists is lower than that of non-altruists, and relative fitness is what determines whether a trait will evolve. Niche-constructing activities which satisfy the criterion $\varepsilon > 0$ but not $\alpha_1 > 1 > \alpha_2$ are analogous to weakly altruistic actions – they boost the absolute fitness of every organism in the population, by causing the resource **R** to accumulate, but the relative fitness of niche-constructors is lower than that of non-niche-constructors. And just as weakly altruistic actions cannot evolve by selection in an unstructured population, even though they boost absolute fitness, so niche-constructing activities which do not confer a relative fitness advantage on their performers will not evolve, despite boosting absolute fitness. The ambiguity between the weak and strong notions of altruism has often caused confusion in the literature on altruism and group selection; it would be a pity if similar confusion were generated by the ambiguity between the relative and absolute notions of 'positive' niche-construction.

Turning to more general matters, one striking feature of **OLF**'s discussion is their liberal use of 'informational' and 'semantic' vocabulary to describe both natural selection and niche-construction, which permeates the whole book. The orthodox Darwinian process of natural selection leading to changes in the genetic composition of a population is described by **OLF** as the accrual of "semantic information", encoded in the DNA, which "expresses the adaptations of organisms" (p. 420). When niche-constructing organisms modify the selective environment, they are "expressing semantic information"; this in turn can lead to the "acquisition of further semantic information" as the modified environment selects for new genes (in either the same or a different population). The net result is that the combined processes of natural selection and niche-construction lead to the "flow of semantic information" through ecosystems, which complements the better known flows of energy and matter that ecologists traditionally describe (pp. 321–335).

The image of information 'flowing' through ecosystems is striking, but **OLF** do not say a great deal about how the concept of semantic information is to be understood, nor about the rationale for using information-theoretic vocabulary to describe natural selection in the first place. A number of biologists and philosophers of biology have explored this issue in recent years, though without reaching a clear consensus (Maynard Smith 2000; Sarkar 2000; Sterelny 2000; Godfrey-Smith 2000; Moss 2003). One source of opposition to the idea that genes contain "information" about organisms' adaptations, as **OLF** hold, is that the causal pathway between DNA and organismic phenotype is extremely complicated. On this view, we can at best speak of genes containing information about, or "coding for", primary protein structure, but not about whole-organism phenotypes. To think otherwise is to ignore the complexity of ontogenetic development, and the fact the genes are but one of many causal ingredients in development. Interestingly, **OLF** come close to making this point themselves, noting that "genes may be regarded as determining proteins, but their influence on phenotypes is more diffuse" (p. 372). However, they apparently regard this as quite compatible with talking about "information, typically encoded in DNA, that specifies the adaptations of organisms. It pertains to the life requirements of organisms in their local environments" (p. 420). I suspect that those sceptical of "genetic information" and kindred notions would regard the diffuseness of the gene/phenotype link, which **OLF** acknowledge, as reason for not describing natural selection and niche-construction in so overtly semantical terms.

Still on the topic of information, the discussion of "information-gaining ontogenetic processes" in section 6.3.2 reveals an interesting theoretical commitment. **OLF** note that many species have evolved "complicated processes that allow individual organisms to acquire *other kinds of information*", in addition to the "genetic information" contained in the DNA (p. 255, my emphasis). They cite as examples the vertebrate immune system and brain-based learning in animals, both of which allow organisms to fine-tune their behaviour to the local circumstances. So **OLF** clearly hold that the sense of "information" in which genes contain information about the organismic phenotype is the same as the sense in which animals can acquire information about their environment by learning, or can acquire information about which antibodies to produce. I have no quarrel with the idea that the ontogenetic processes in question, particularly animal learning, can be described as a type of information acquisition. But it is open to debate whether this is "information" in the same sense as that in which genes carry "information", if indeed there is a respectable sense of the latter.

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? One possible reason is this. Suppose one follows Dennett (1995), Millikan (1984) and others in regarding natural selection as the ultimate source of all "meaning" or "intentionality" in the world. The rationale for talking about genetic information, on this view, is that genes are the outcome of a long process of natural selection, which has invested them with "meaning". I suspect that **OLF** would be happy with such a view. Suppose one also holds that ontogenetic processes such as brain-based learning, antibody production in the vertebrate immune system etc, are *themselves* Darwinian processes, that operate over ontogenetic rather than phylogenetic time. (There is of course a long tradition of thinking about the immune system in just this way.) Then, one will naturally be led to describe these ontogenetic processes in informational terms too, thus securing the analogy with genetic information.

I was unsure if **OLF**'s reason for regarding DNA-encoded information and ontogenetically-acquired information as two types of the same basic stuff, "information", stemmed from this line of reasoning or not. They certainly express sympathy with the idea that genetic information is properly so-called in virtue of natural selection having shaped phenotypes, and the genes which specify them, over a long period of time. They also say that ontogenetic processes such as learning and the immune response "can be regarded as operating in a manner loosely analogous to the Darwinian algorithm" (p. 256). However, they then go on to point out salient *differences* between brain-based learning and Darwinian natural selection (pp. 256–257). In my view these differences are salient enough to undermine the analogy between animal learning and natural selection altogether; if anything, learning is better modelled as a Lamarckian rather than a Darwinian process. (The immune system provides a far better analogy with Darwinian selection.) If this is right, but we still wish to regard animal learning as a type of information acquisition, as we surely should, it follows that the link between natural selection and information acquisition cannot be so intimate after all. Non-Darwinian processes can lead to information acquisition too. But then one is left wondering about the rationale for talking about genetic information, and about whether there is a univocal sense of "information" in which genes contain it and ontogenetic processes permit it to be acquired.

Finally and possibly relatedly, I did not quite understand why **OLF** wished to describe niche-construction as a *selective* process. In chapter 4 they tell us that evolution "depends on two selective processes rather than one: a blind process based on the natural selection of diverse organisms in populations exposed to environmental selection pressures, and a second process based on the semantically informed selection of diverse actions, relative to diverse environmental factors, at diverse times and places, by individual niche-constructing organisms" (p. 185). When a bird builds a nest, or a spider spins a web, there is perhaps a sense in which they are engaging in a selective process. Not just any twigs will make a good nest, and not just any location will make for a good web, so the bird must 'select' appropriately sized twigs and the spider must 'select' an appropriate location for its web. But this seems to me a relatively trivial sense of 'selection', and not relevantly similar to the 'selection' in Darwinian selection. Indeed in this broad sense of 'selection' virtually any activity can be characterised as selective. I think **OLF** are quite right to insist that niche-construction should be regarded as part of the process of evolution, rather than merely a by-product, and that it therefore complements natural selection as one the forces driving organic evolution. But this does not compel us to regard niche-construction as a *selective* process, for not all evolutionary processes need be selective.

In conclusion, I should say that *Niche Construction* seems to me a marvellous achievement, worthy of close study by anyone with an interest in evolution and ecology. **OLF** present a sustained, rigorous, and highly original argument for the extended evolutionary theory they advocate, that blends theoretical, empirical and philosophical considerations in a most impressive way. Moreover, unlike certain authors who have called for a major re-orientation or re-structuring of evolutionary theory, such as Gould (2002), **OLF** provide positive and practical suggestions for how other researchers can put their evolutionary ideas into practice. The book merits a wide audience among biologists and philosophers of biology alike.

References

Dennett, D.: 1995, *Darwin's Dangerous Idea*, Penguin, London.

Godfrey-Smith, P.: 2000, 'Information, Arbitrariness and Selection: Comments on Maynard Smith', *Philosophy of Science* **67**(2), 202–207.

Gould, S.J.: 2002, *The Structure of Evolutionary Theory*, Harvard University Press, Cambridge MA.

Kerr, B. and Godfrey-Smith, P.: 2002, 'Individualist and Multi-level Perspectives on Selection in Structured Populations', *Biology and Philosophy* **17**, 477–517.

Maynard Smith, J.: 2000, 'The Concept of Information in Biology', *Philosophy of Science* **67**(2), 177–194.

Millikan, R.: 1984, *Langauge, Thought and other Biological Categories*, MIT Press, Cambridge MA.

Moss, L.: 2003, *What Genes Can't Do*, MIT Press, Cambridge MA.

Sarkar, S.: 2000, 'Information in Genetics and Developmental Biology: Comments on Maynard Smith', *Philosophy of Science* **67**(2), 208–213.

Sterelny, K.: 2000, 'The "Genetic Program" Program: A Commentary on Maynard Smith on Information in Biology', *Philosophy of Science* **67**(2), 195–201.

Wilson, D.S.: 1980, *The Natural Selection of Populations and Communities*, Benjamin/ Cummings, Menlo Park, CA.