Biol Philos (2008) 23:1–25
DOI 10.1007/s10539-007-9061-2

Biological Levers and Extended Adaptationism

Gillian Barker

Received: 19 October 2006 / Accepted: 30 January 2007 / Published online: 20 March 2007

Abstract Two critiques of simple adaptationism are distinguished: anti-adaptationism and extended adaptationism. Adaptationists and anti-adaptationists share the presumption that an evolutionary explanation should identify the dominant simple cause of the evolutionary outcome to be explained. A consideration of extended-adaptationist models such as coevolution, niche construction and extended phenotypes reveals the inappropriateness of this presumption in explaining the evolution of certain important kinds of features—those that play particular roles in the regulation of organic processes, especially behavior. These biological or behavioral 'levers' are distinctively available for adaptation and exaptation by their possessors and for co-optation by other organisms. As a result they are likely to result from a distinctive and complex type of evolutionary process that conforms neither to simple adaptationist nor to anti-adaptationist styles of explanation. Many of the human features whose evolutionary explanation is most controversial belong to this category, including the female orgasm.

Keywords Adaptation · Adaptationism · Exaptation · Regulation · Feedback · Coevolution · Niche construction · Extended phenotype · Evolutionary explanation

Explaining why

> Walking one evening along a deserted road, Mulla Nasrudin saw a troop of horsemen coming towards him. His imagination started to work; he saw himself captured and sold as a slave, or impressed into the army. Nasrudin bolted, climbed a wall into a graveyard, and lay down in an open tomb. Puzzled at his strange behavior, the men – honest travelers – followed him. They found him stretched out, tense and quivering.

G. Barker (✉)
Department of Philosophy, Bucknell University, Lewisburg, PA 17837, USA
e-mail: gbarker@bucknell.edu

> “Why are you down in that grave? We saw you run away. Can we help you?” “Just because you can ask a question does not mean there is a straightforward answer to it,” said the Mulla, who now realized what had happened. “It all depends upon your viewpoint. If you must know, however: *I* am here because of *you*, and *you* are here because of *me*” (adapted from Shah 1972, p. 16).

Many of those in the business of evolutionary explanation resemble the bewildered travelers in this Sufi tale. Confronted with a state of affairs that cries out for explanation, they seek what the Mulla calls a ‘straightforward answer’ to the question, ‘Why are things this way?’—an answer that identifies a simple cause of the phenomenon to be explained. Elizabeth Lloyd’s (2005) recent examination of debate on the evolution of the female orgasm offers a fascinating look at straightforward evolutionary explanation-making in action. Most participants assume that the female orgasm is an adaptation for *something*, but their views about *what* it is an adaptation for are wildly disparate: the adaptive roles assigned to the female orgasm range from reinforcing the pair bond to assisting in fertilization, from guiding females’ choice of mates to selecting superior spermatozoa. Lloyd shows that none of the adaptationist explanations that have been offered is currently well-supported empirically, and argues that the evidence best supports the hypothesis that the female orgasm is not itself an adaptation but a developmental by-product of the adaptive male orgasm (Symons 1979). All of the hypotheses that Lloyd canvases, including the non-adaptive by-product account, represent attempts to find a straightforward answer to the question ‘Why?’. I’ll argue here that examining the relationship between two critiques of adaptive explanation reveals that seeking a ‘straightforward answer’ is inappropriate for certain important kinds of cases—including the female orgasm—and suggests resources for developing an approach to evolutionary explanation capable of accommodating such cases. At the end of the paper I will consider what this approach suggests about the evolution of the female orgasm. A good deal of preliminary work must be done in order to shed light on this controversial and, as it will turn out, very complex case.

Anti-adaptationism and extended adaptationism

Two lines of critique have been prominent in recent discussions of adaptive explanation. The *anti-adaptationist* critique says essentially that those who offer adaptive explanations are too quick to suppose that the feature to be explained is an adaptation at all. What I’ll term the *extended adaptationist* critique takes issue instead with the particular notion of adaptation employed in most adaptive explanations. Anti-adaptationism and extended adaptationism are not competing positions, but strands of thought that are often closely entangled; noticing the differences between them has far-reaching consequences for our understanding of adaptation, however. I shall argue that the influential anti-adaptationist critique requires substantial revision in light of some of the implications of extended adaptationism. Certain features, because of the kind of role they play in the causal structure of the organism and in its interactions with other organisms, are *apt for aptation*—for adaptation, exaptation and co-optation. For these features the adaptationists’ presumption that the explanation of their presence must be an adaptive story is approximately correct. For the same reasons, however, the sort of adaptive story required is likely to be much more complex than the notorious ‘just-so stories’ of simple adaptationism.

Anti-adaptationism

Anti-adaptationists, in a critique epitomized by Gould and Lewontin’s ‘‘Spandrels of San Marco’’ (1979) and Gould’s early contributions to the debate on the female orgasm (Gould 1987; Lloyd 2005), attack the error of *adaptationism*: the uncritical presumption that all features of an organism are effectively-independent traits produced by adaptive evolution, so that natural selection is the only significant force driving evolutionary change. They argue that before asking the adaptationists’ question, ‘What is this feature an adaptation for?’, we should always first ask the question ‘Is this feature an adaptation at all?’ and warn that an affirmative answer to this question requires careful support including direct evidence of an appropriate history of selection.

The anti-adaptationist critique identifies adaptationism with two basic tenets: the view that the features of an organism are effectively discrete and mutually independent (*trait atomism*) and the view that each feature is an adaptation, optimally-designed for its function by natural selection within the constraints imposed by trade-offs among conflicting adaptive demands (*panselectionism*) (Gould and Lewontin 1979; Gould and Vrba 1982). From these basic presumptions the research practices characteristic of adaptationism are taken to follow, including the practices of inferring a feature’s adaptive history from its present utility and of trying one adaptive explanation after another rather than considering the possibility of non-adaptive explanations. Critics of adaptationism argue that the presumptions of trait atomism and panselectionism are demonstrably false, and that the practices based on them should therefore be abandoned (Gould and Lewontin 1979; Lloyd 2005). Defenders of adaptationism reply that the presumptions should be accepted as useful approximations to the truth, as illuminating the important parts of evolution, or as methodologically valuable in some other way, which is to say precisely that the practices based on them are worth pursuing (Mayr 1983; Dawkins 1986; Dennett 1995; see also Godfrey-Smith 2001).

Now, nobody really believes that the principles of trait atomism and panselectionism are literally true, or even that the practices associated with them are universally useful as methodological principles. Even those who regard themselves as ‘ardent adaptationists’ (Alcock 1987; Barash 2005) do not think that it is true (or useful to assume) that literally *every* present utility corresponds to an evolved adaptation. On the other hand, nobody really adheres to the stringent standards of strict anti-adaptationism either, believing that it is *never* legitimate to infer from present utility to adaptive history. It appears, in fact, that both adaptationists and anti-adaptationists tacitly accept that some features are good candidates for adaptive explanation while others are not. Adaptationists and anti-adaptationists clearly disagree about which features, and how many, belong in which of these two categories, but there has been no explicit examination of the principles of classification at work here. *Are* there in fact *kinds* of features (distinguishable by characteristics they possess at present, rather than their histories) that are especially likely or unlikely to be adaptations? If so, there may be *specifiable conditions* under which the methodology of adaptationism is appropriate.

Adaptationism and anti-adaptationism do not provide obvious grounds for thus dividing features into likely and unlikely candidates for adaptive explanation. Instead, each approach treats all features alike. For any feature, the adaptationist rule is to look first for an adaptive explanation, while the anti-adaptationist rule is to consider the full range of explanatory alternatives including selective adaptation, drift, allometry, developmental constraints, and so on, and in particular *not* to accept a speculative adaptive explanation without the relevant sort of historical support. On an atomistic view of organisms this sort of one-size-fits-all policy seems appropriate. But anti-adaptationists reject the adaptationist's atomism—Gould and Lewontin (1979, p. 583) emphasize that ''organisms are integrated wholes, not collections of discrete objects''—and holism implies the possibility of deep differences among kinds of features, resulting from differences in their degree and kind of causal integration with other features of the organism. Anti-adaptationists have indeed noted the implications, for evolutionary explanation, of the ways that features are connected through the mechanisms of inheritance (e.g. linkage and pleiotropy) and development (e.g. allometry and developmental byproducts). They have paid little attention, however, to interconnections among features in the current causal structure of the living organism. The important implications of these interconnections become plain, however, when we consider some aspects of the extended adaptationist critique of adaptationism.

Extended adaptationism

The second line of critique of standard adaptationist explanation appears in a diffuse and various literature that takes issue with the particular notion of adaptation used in most adaptive explanations. According to a view shared by many adaptationists and anti-adaptationists (Lloyd 2005, p. 157), adaptation is the counterpart of evolutionary function, so that a feature is an adaptation for producing a particular effect if (and only if) that effect is its evolutionary function, i.e. if the feature is the result of past selection for the production of that effect. The selective process that produces adaptive evolution is commonly understood as involving a population of *individual organisms* with a *genetic system of inheritance*, acted upon by an independent *selecting environment*. Each of the interrelated elements in this characterization has been called into question, however. Some writers argue that ordinary adaptive explanations are mistaken in attributing adaptations only to individual organisms when in fact adaptive processes involving other sorts of biological objects may play important roles in evolution (Buss 1987; Turner 2000), offering adaptive accounts involving selection acting on social groups (Sober and Wilson 1998), on developmental systems (Oyama 2000; Griffiths and Gray 1994; Oyama et al. 2001), or on genes via their extended phenotypes (Dawkins 1982). Some note that there may be important and pervasive non-genetic systems of inheritance in organisms or in other units of selection, such as cytoplasmic inheritance, maternal effects, cultural inheritance and the inheritance of modified environments (Jablonka and Lamb 1995; Jablonka 2001; Boyd and Richerson 1985; Cavalli-Sforza and Feldman 1981; Wolf et al. 1998).

Some point out that a bifurcation of the world into selecting environment and population under selection is misleading, since organisms exert selective pressures on one another (Ehrlich and Raven 1964; Janzen 1980; Futuyma and Slatkin 1983;

Thompson 1995). Others argue further that the whole notion that independent environments exert selective pressure on populations of passive organisms—that the genes propose and the environment disposes—is deeply mistaken; that, on the contrary, organisms are constantly and actively remaking their environments to such an extent that it is sometimes more accurate to say that the environment is adapted to the organisms than the reverse (Lewontin 1983a, b, 1984; Bateson 1988; Laland 2004; Laland et al. 1996, 2000; Day et al. 2003; Odling-Smee et al. 2003).

These revisions of the basic evolutionary account of adaptation are diverse; some of their authors regard themselves as critics of adaptationism, others regard themselves as refining and extending a good idea. A common thread nonetheless unites these various approaches: each is somehow animated by the idea that before we get to the adaptationists' question, 'What is this feature an adaptation for?' we must first resolve the question '*What* is adapting to *what*?' (Lewontin 1983b), and each treats adaptation as a process that somehow extends beyond the boundaries of the individual organism, so that '*what* is adapting' may be something larger than an individual organism, such as an extended phenotype, a social group, or an organism–environment complex. With appropriate respect for the real differences among these approaches and for the rich ambiguity of Lewontin's question (What kinds of objects and environments are we talking about as being adaptively related? Which way does the adaptive arrow connecting them go?) I'll refer to this sort of approach generally as *extended adaptationism.*

Commonalities and conflicts

Many critics of simple adaptive explanations regard both anti-adaptationism and (some) extended adaptationism as well-motivated and think of them as similar or mutually supporting. There is surely something right about this perspective. Both anti-adaptationists and extended adaptationists take seriously the idea that attributions of adaptation are essentially explanatory, that to say what a feature (of some biological object or system) is *adapted for* is to explain not just what it does or even what it is useful for, but *why it is present*—how it came to be or at least how it has been maintained. On the other hand, both recognize that simple adaptive explanations often oversimplify the complex and various factors that interact in the processes involved in producing evolutionary change and stability. Anti-adaptationists point out that there is more to evolution than adaptation, and extended adaptationists point out that there is more to adaptation than simple individual selection. Yet in their details the two approaches are not easy to reconcile.

The anti-adaptationist approach focuses on distinguishing those features of organisms that are correctly explained as adaptations from those that are not. This focus has some notable effects. In their effort to counteract uncritical adaptationism, anti-adaptationists often set restrictive standards for what is to count as an adaptation, accepting the adaptationists' simple model of what adaptation is and then pointing out that many features that seem adaptive in the sense of being useful to the organisms possessing them nonetheless do not have, or cannot be shown to have, the kind of evolutionary history required for the attribution of adaptation. Like adaptationists, many anti-adaptationists see adaptations as the solutions produced by natural selection to pre-existing problems set by the environment. As a result, they tend to treat adaptation as an accomplishment rather than as a process, to ask of any feature whether it *is* an adaptation (a solution to an environmental problem) rather than what role adaptation *played* in its evolution. Many extended adaptationists see adaptation instead as the ongoing process of *adapting*, a process that is never completed because the organisms (or other systems under selection) are constantly acting on their environments in ways that change the selective forces that act upon them, creating new problems with every solution. Thus, where the anti-adaptationist sees a sharp distinction between simple selective adaptations and other—non-adaptive—features of organisms, the extended adaptationist sees a diverse array of evolutionary processes, in many of which adaptation and other factors interact in complex ways, and many of which involve forms of adaptation that are quite different from simple selective adaptation.

The central insights of anti-adaptationism and extended adaptationism are not incompatible, but the tensions noted above leave us with some questions in need of resolution. How common, and how important, are evolutionary processes that are neither simple problem-solving adaptation nor strictly non-adaptive? Are there importantly different types of extended adaptation to be distinguished? How ought we to proceed methodologically—under what circumstances ought we to start with the expectation of some sort of adaptive explanation, and under what circumstances ought we to resist such an expectation?

This paper will not resolve all of these questions, but it will give some indication of the shape of the territory that they open up and discern some paths that are worth following. The extended adaptationist perspective reveals the possibility of deep distinctions among kinds of features resulting from their differing causal roles in the complex processes of evolution, and offers resources for identifying at least one important category of features that are good candidates for extended adaptive explanation and that are therefore appropriate cases for the application of a modified form of methodological adaptationism. These features have a special evolutionary status because of the kind of role they play in regulating organisms' functioning and behavior and hence in mediating the influence that organisms have on their own selective milieus and on those of other organisms in what I'll call "selective interaction."

Selective interaction: coevolution, niche construction and extended phenotypes

Organisms modify their environments in ways that change the selective pressures acting upon them. Furthermore, an overwhelmingly important portion of every organism's environment is composed of other organisms, so what each organism does—what effects it has on the world around it—has selective effects on other organisms as well. These facts about the structure of the organic world are glaringly obvious, yet their diverse manifestations and far-reaching implications are difficult to capture adequately in adaptive explanation. I'll consider three extended-adaptationist approaches to such selective interactions that are variously useful in thinking about their implications.

Coevolution

Darwin recognized that interactions among organisms can result in reciprocal influences on selection both within species (especially in the form of sexual selection) and between species. Ehrlich and Raven (1964) termed the special form of adaptive evolution that results from such interactions *coevolution*. Paradigm cases of coevolution are those involving two species interacting in such a way that each exerts selective pressures on the other and the resulting selective processes form causal feedback loops. Thus, for example, plants and their insect pollinators may evolve highly specialized mutualistic relationships through a process of coevolution in which changes in flower structure select for changes in pollinator behavior, which in turn select for changes in flower structure, and so on. Coevolution is often regarded primarily as a mechanism that explains the evolution of mutualism or symbiosis, but it is often antagonistic instead, as in the case of 'arms races' between predators and prey or hosts and parasites. Coevolution may also occur within a single species: sexual selection, feedback between selection for mechanisms for signal emission and signal reception, and intraspecies arms races are all examples of intraspecies coevolution.

Niche construction

The coevolution model provides a means of understanding relatively direct and tightly-reciprocal selective interactions, but organisms influence each other's selective environment in many ways that are less direct or more diffuse but may nonetheless be evolutionarily important. The diverse ways in which organisms change their environments and thus change the selective forces acting upon them and upon each other—dubbed ''niche construction''—have been explored recently by Odling-Smee (Odling-Smee et al. 2003; Laland 2004; Laland et al. 1996, 2000) and others.

Niche-construction theorists emphasize that organisms are not merely the passive objects upon which the selective forces of their environments act; they actively change their own environments in a wide variety of ways. The processes by which organisms change their own selective environments are both diverse and pervasive. Two basic categories of niche construction are distinguished: 'perturbational' and 'relocatory' (Odling-Smee et al. 2003, p. 44). In perturbational niche construction, organisms actually modify the properties of their surroundings. The literal construction of nests, webs and burrows and the storage of food or other resources are obvious forms of perturbational niche construction. But more subtle effects on the environment are also included: the release of esters, pheromones or allopathic chemicals; the emission of sounds, even casting shade or creating wind shelter are forms of niche construction. All organisms consume resources and excrete waste products, and these processes may change their local environments considerably. Such processes involving large populations can have massive cumulative effects on the macro-environment: the role of earthworms in the creation of topsoil, of plants in the creation of earth's current oxygen-rich atmosphere, and of forests in creating large-scale weather patterns are all obvious examples. In relocatory niche construction, organisms 'change' the environment that acts upon them by affecting what part of the larger environment they interact with. Animal migration, choice of nesting site, and general habitat selection are obvious forms of relocatory niche construction. Plants less obviously affect the micro-environments that they or their offspring inhabit by various mechanisms of dispersal or growth.

Niche construction is thus a process that can establish a feedback relationship between the niche-constructing organisms and their selective environment, and so in turn between two sorts of traits in the niche-constructing organisms: traits that contribute to niche-construction, and traits that are selectively affected by the effects of niche construction. Niche construction can have effects on the selective forces impinging on many organisms other than the niche-constructing species, resulting in complex and diffuse reciprocal relations among the selection pressures on many traits in many different organisms. These effects are most obvious in cases of large-scale perturbational niche construction, but can also result from habitat selection—this indeed is the principle reason for treating the latter as a form of niche construction.

The extended phenotype

The niche-construction perspective has important commonalities with that articulated by Richard Dawkins in *The Extended Phenotype* (1982). On this view, *all* effects that a gene has on its surroundings—whether inside or outside the boundaries of the organism—are considered part of the gene's phenotype, *provided that they have the potential to play a part in determining the gene's replicative success*. Thus a bird's nest or a spider's web is part of the extended phenotype of each of the bird genes or spider genes involved in producing them. As Dawkins emphasizes, a gene's extended phenotype extends not only beyond the boundaries of the organism containing the gene, but across the boundaries of other organisms. If a gene in one organism has effects that result in changes in the structure or activities of another organism, those effects may be part of the gene's extended phenotype. Thus the structures produced by host organisms under the influence of parasites, such as the plant galls induced by wasps, are parts of the extended phenotypes of the parasites' genes. Behavior of one animal that is provoked by a signal emitted by another is part of the extended phenotype of the signaling organism's genes. Like the niche-construction perspective, Dawkins' view recognizes that organisms do change their environments, including other organisms, in ways that in turn affect the selective forces that act upon them.

Key elements of selective interaction

Each of these three models captures some important aspects of selective interaction among organisms, and each overlooks or at least obscures some such aspects.

The key elements for my purposes are as follows:

(i) Very important parts of each organism's environment are made up of other organisms, both conspecifics and heterospecifics.
(ii) Organisms actively modify both biotic and abiotic elements of their environments and actively select which parts of the larger environment they inhabit or interact with.
(iii) These activities often have fitness consequences for the organisms that perform them, for any organisms directly acted upon, and for other organisms.

(iv) Because each organism is both a cause and a recipient of fitness consequences, the above activities can result in feedback loops with typical features such as the capacity for stabilization (via negative feedback) or self-reinforcing change (via positive feedback).

Each of the three extended-adaptationist models sketched above develops or obscures elements (i)–(iv). The coevolution model gives an illuminating emphasis to the direct selective effects that organisms have on one another, and on the powerful feedback processes that can result, but does not easily take account of the many equally important indirect ways that organisms can affect their own and others' selective milieus. The niche construction model provides a representation of organisms' selective effects and their capacity to produce feedback processes that is potentially extremely general, but as developed to date (Odling-Smee et al. 2003) it tends to overemphasize literal constructions like nests and webs and to downplay such ephemeral effects as signal emission. Further, because a niche is typically conceived as belonging to a species rather than an individual, it is difficult to include direct effects that organisms have on conspecifics as forms of niche-construction. The extended-phenotype model gives a rich representation of the many ways that an organism's genes can affect their own chances of replication, via both long-term and ephemeral effects on other organisms and on the abiotic environment, but does not take account directly either of the effects that one gene has on the fitness of other genes (in the same organism or in another) or of the resultant feedback processes.

Feedback

The concept of causal feedback, we have seen, plays an important role in many extended-adaptationist approaches to evolutionary explanation. The theory of coevolution is essentially an elaboration of the idea that two populations can be related in such a way that the selective effects they exert on one another form a feedback loop: coevolution in the narrow sense has occurred when ‘‘a trait of one species has evolved in response to a trait in another species, which trait itself has evolved in response to the trait in the first’’ (Futuyma and Slatkin 1983, p. 1). Niche construction theory emphasizes the role that feedback between organisms and their environment plays in shaping evolution:

> The evolutionary significance of niche construction hangs primarily on the feedback that it generates.... [F]eedback loops introduce major instabilities, associated with positive feedback cycles, and hyper stabilities, associated with negative feedback cycles. Feedback can produce ‘lock-in’ effects,...timelags in the response to selection, and evolutionary momentum. (Laland et al. 2000; see also Robertson 1991).

The notion of feedback is of course not new to biology. Its more familiar use is in theories of biological regulation deriving originally from the work of Claude Bernard (1865) on homeostasis and of Norbert Wiener (1961) on cybernetics or theory of control in engineered and biological systems (Ashby 1956; Cziko 2000; Keller 2002). The possibility of regulation or control results in part from the same peculiar features of feedback structures that interest the proponents of niche construction: their capacity to amplify or suppress change. The three extended-adaptationist approaches that I've described do not explicitly address the issue of regulation. Nonetheless, as we shall see, regulation has a role of the first importance to play in extended adaptation, through the interaction between the feedback structures involved in regulation and the feedback structures involved in selection. To the list of key elements of selective interaction in the previous section, therefore, we must add the following:

(v) Some of the activities discussed in (i)–(iv) are regulated.

Regulation

Before turning to the implications of regulation for selective interaction, it will be useful to develop a simple general picture of biological regulation. Organisms are regulative systems—they actively maintain or systematically adjust their internal states with respect to many different variables, against backgrounds of changing external conditions and in ways that are tuned to changing conditions both internal and external to the organism. The most familiar form of regulation is the sort of simple negative feedback control process involved in many biological homeostatic systems and in artificial cybernetic or control systems (Wiener 1961; Ashby 1956). Recent studies of dynamical systems explore more complex mechanisms of regulation, but many important processes of biological regulation are usefully modeled by means of the simple negative feedback control process, and many of the lessons that this sort of regulation has for thinking about selection and adaptation apply equally to more complex forms of regulation. The details of simple regulation will be important for the argument that follows, so I shall review this familiar sort of control structure in some detail.

A negative feedback control process works to regulate some variable, keeping it close to a certain predetermined value by counteracting any disturbance that moves the regulated variable away from that value. To do this requires mechanisms for sensing the current state of the regulated variable, comparing it to the predetermined reference value, and causing the regulated variable to change so as to reduce the difference between the two. These include:

- A *reference* that fixes the predetermined 'goal' value for the regulated variable (the reference value may itself vary according to conditions outside the control process, but in the simplest case it is constant.)
- An *indicator* whose state is determined by (and which thus 'senses') the present value of the regulated variable.
- A *feedback loop* that communicates the state of the indicator to the comparator.
- A *comparator* that compares the state of the indicator with the reference. If a difference is detected, the comparator sends an *error signal* to the effector.
- An *effector* that is induced by the error signal to modify the regulated variable in such a way as to reduce the error (the difference between reference and indicator values).

In the familiar case of a home heating system, for example, the reference value is the temperature you set the thermostat to maintain; the indicator is the thermometer that registers the room temperature; the comparator is the mechanism in the thermostat that registers whether the thermometer reading is within or outside the acceptable range around the reference value, and the effector is the furnace that acts to bring the indicator state back into the acceptable range. This is a feedback or ‘closed loop’ system because the value of the regulated variable and the state of the effector are reciprocally related in a closed causal loop, so that a change in the regulated variable causes a change in the effector, which in turn causes a change in the variable. It is a *negative* feedback system because change in one direction in the regulated variable has further effects (here, via the effector) that then produce change in the opposite direction in the original variable. It’s important to note, however, that there is more to a negative feedback control process than negative feedback alone. Simple negative feedback between two variables by itself can produce local stability, as in the familiar case of dynamic equilibrium in population densities of predators and prey. The number of hares increases, and the number of foxes rises as a result. The now numerous foxes eat more hares, so the hare population drops, causing the fox population to decline as well. As the number of foxes falls, the hare population begins to rebound, and the cycle continues. The result is that both populations are maintained at a relatively stable level if outside conditions do not change radically. But this is not yet a case of regulation. Regulation allows a system to maintain stability against a changing background of environmental disturbances, as in the case of thermoregulation in mammals (which may maintain a stable internal temperature despite extreme variation in external temperature) or to track a changing reference condition, as in the case of heart rate regulation (which allows the heart rate to vary to meet changing demands). These capacities require the special features of the control process: the reference, indicator, and comparator that enable the system to evaluate and correct the relationship between its current state and the reference value.

More complex control systems may comprise hierarchies or webs of component systems. In a hierarchy, the elements of higher-level systems (such as the effector) themselves contain lower-level control systems, whose reference values can therefore vary in accord with the demands of the higher-level system. In a web, a collection of interrelated systems may be connected so that the mechanisms that each employs call on the capacities of others. Organisms effectively regulate vast arrays of interrelated variables, from those that constitute their internal states, such as temperature, water balance, ion balance and pH, to external variables such as the state of their local environment. They do this by means of both internal processes and interactions with the environment, including behavior, and many of the processes that thus serve as components of regulatory systems are themselves regulated.

Animal behavior serves as part of the effector for many different control systems within the organism; it is able to contribute as effectively and flexibly as it does to the working of the organism because it is itself regulated. In the regulation of animal behavior the effectors are the mechanisms that enable the organism actually to move around and do things, comprising the muscles that produce motion together with the neural motor system that immediately directs their activity. The indicators are sensations, internal and external, that register the current state of the organism. References tell the organism what indicator state to aim for in a given context, and include both innate and learned stimulus–response patterns, preferences, attractions, aversions and so on. Error signals register how far the organism is from the reference state, and induce behavior accordingly; in some animals they are associated with feelings of pleasure and discomfort or dissatisfaction. These feelings themselves are not required for immediate regulation of behavior—to be effective an error signal need not produce any particular subjective experience for the organism—but they do play a crucial role in learning, i.e. as mechanisms for adjusting reference states. While organisms do a great deal of regulation of internal and external processes that does not involve behavior, the regulation of behavior has certain important implications for evolutionary explanation stemming from the special role that behavior plays in mediating organisms' interactions with their environments (Bateson 1988; Plotkin 1988). These implications make behavioral regulation worth exploring more closely. Some of the lessons that we learn from behavioral regulation will also turn out to have wider application.

Biological leverage

Any biological system that regulates some variable, as we have seen, contains a mechanism capable of substantial causal action, the effector, which can be set in motion by changing the state of either of two components that are highly responsive to impinging causal influences, the indicator and the reference. The indicator must be responsive to causal influences, since its function is precisely to track changes in the conditions internal or external to the organism. The reference that fixes the goal-state in a control process that is part of a complex hierarchy or web of negative control processes will typically also be relatively easy to modify, so that it can be adjusted according to the demands of the other processes in the larger system. What this means is that a control system can effectively act as something analogous to a *lever*: a causal structure that transforms a small initial cause (acting on the reference or the indicator) into a much larger effect. A small cause acting on either the reference or the indicator of my home thermostat produces a large effect in the activity of my furnace. A small change in the goal states or sensory perceptions of an organism can produce a large effect in its behavior. Such systems present obvious advantages for their possessors. If it is beneficial to an organism to change what an effector is doing, the reference provides an easy and safe means of doing this; easy because the reference is designed to be responsive to the causal processes by which it can be modified by selection, learning, or adjustment to meet the demands of other control processes in the web, and safe because the reference has influence over its own control system alone while the effectors are usually involved in many other regulative systems so that the risk of detrimental side-effects to modification of the effectors is much greater. It is easier and safer to change the temperature in my house by changing the thermostat setting than by modifying the furnace itself, and easier and safer for an organism or a lineage of organisms to change its rate of sugar intake by changing its flavor preferences than by modifying its mouth structure or digestive apparatus. By the same token, however, such systems are inherently vulnerable to exploitation by other organisms. What the effector does depends on the state of the indicator, and in order to serve its function the indicator must be both accessible and responsive to outside intervention. If an organism uses a preference for sweet foods as a reference for regulating its eating behavior, other organisms can modify that organism's behavior by presenting it with sweet-tasting substances that act as lures or rewards.

An organism can change the reference involved in the regulation of one of its own behaviors in a number of different ways. It may possess a set of alternative values that can be temporarily activated according to context: a foraging animal, for example, may shift from one search image to another as local conditions change, or an animal's rate of food consumption may change radically when the external temperature drops below a certain level and hibernation begins. It may add a new reference to its repertoire (or lose an old one) by means of learning: the organism acquires a new stimulus–response association, or a new attraction or aversion. If the new reference is different enough from any in its existing repertoire we say that the organism has learned a 'new behavior,' like the birds that learned to use their existing effectors to the novel end of pecking holes in the lids of milk bottles (Fisher and Hinde 1949). References may also undergo evolutionary change, so that a population evolves a new preference or innate stimulus–response pattern, for example. In both individual and evolved changes the modification of the reference often reflects a change in what the regulatory system is called upon to do in contributing to some higher-level process of regulation. Sometimes, however, the modification in the reference reflects a shift in *what* higher-level process the system is contributing to. Such a change may be described as an *exaptation* insofar as it results in the recruitment of an existing effector to a new end. Thus, according to Gould and Vrba (1982), birds' wings may have evolved from a precursor structure that was used to slow or speed loss of body heat, so that an effector (the wing structure and the muscular and neural mechanisms that move it) originally used to regulate body temperature came to be used to regulate aerial locomotion, employing a new repertoire of references to do so.

The structure of a control system thus allows it to contribute flexibly to the economy of the organism, but also makes it vulnerable to exploitation by other organisms. As the extended adaptationist perspective emphasizes, what one organism does often has important consequences for other organisms. The highly responsive indicator in a system that regulates behavior plays a key role in determining what the organism will do, so a small effort invested in affecting an indicator in one organism may have a large payoff for another. This is the kind of interaction that is often described as 'manipulation' of one organism by another. The means by which one organism can thus modify another's behavior are diverse. They include all the features or activities usually classified as forms of signaling as well as some that usually are not: colors, forms, scents or flavors that attract pollinators, mates or prey, or that repel predators or act as camouflage; mimetic coloring or behavior; warning calls and mating displays; and so on. Sometimes the effects produced by these influences are detrimental to the interests of the recipient organism; these are the cases that we typically describe as manipulation, exploitation, or deceit. Sometimes the outcome is beneficial to the recipient organism; we commonly call these cases of mutualism, co-operation or communication. The lack of good words for the use of one creature by another that are free of value-laden implications about the effect on the creature being used creates a real terminological difficulty for attempts to examine the general structure of interactions of this sort. In the remainder of this paper I shall use the term *co-optation* with the intent that it be taken to carry no implication about whether the effect of co-optation is good, bad or neutral for the recipient.[1]

---

[1] Gould and Vrba (1982) use 'co-optation' as a synonym for the new term they are introducing: 'exaptation.'

In both exaptation of a system by its possessor and co-optation by another organism, an existing mechanism is redirected to serve a new regulative end. The difference lies in whose regulative interests are served by the change. In exaptation, the redirected mechanism contributes to the regulation of a new variable within the same organism, whereas in co-optation it contributes ultimately to the regulation of some variable within another organism. The means by which the change is effected may also differ: the most effective way to redirect a control system is usually to change the reference, so in most cases of exaptation the reference is modified. In a case of exaptation cited by Preston (1998), herons spread their wings to shade the water so as to see fish more easily; in the regulation of wing movement, they have acquired a new reference state. But the references of one organism's control systems are not easily accessible to influence by other organisms, so the most efficient means for one organism to redirect another's regulatory mechanisms is commonly to modify the input to the indicator. The plant that attracts pollinators with colored and scented flowers, the angler fish that lures prey within reach with a worm-like appendage, and the cat that intimidates attackers by increasing its apparent size all are modifying the behavior of other organisms by affecting the other organisms' sensory input in a way that makes use of existing control processes. Where the co-opting organism has access to some of the recipient organism's internal processes, as in the case of endoparasites, it may be able to modify references for behavioral regulation systems (Thomas et al. 2005).

Another important difference between exaptation and co-optation is in the scale of the potential benefit involved: exaptation may be an efficient means for an organism to redirect its own activities and so to optimize its use of its own resources, but co-optation is a means of tapping into the resources of another organism at very low cost—a potentially tremendous benefit. Financial leverage allows investors to create a return on someone else's money; biological leverage allows organisms to benefit from someone else's capacities and resources. In both cases the 'someone else' may be helped, harmed or left unaffected by the relationship, but often the effect is powerful.

Leverage and selection

The co-optation of the systems of one organism to serve the ends of another can produce powerful fitness effects on both participants. I'll call the organism doing the co-opting the ''agent'' and the organism whose process is co-opted the ''recipient.'' As we have noted, co-optation does not always impose a fitness cost on the recipient organism; it can instead result in a mutualistic relationship that benefits both participants. In a co-optation that is beneficial to both agent and recipient, the intervention of the agent takes advantage of, but does not change, the relationship between the indicator state and the state that the indicator registers. Such cases can be identified with what signal theorists call 'honest signals' (Maynard Smith and Harper 1995), and can arise where the interests of agent and recipient organism are aligned at least up to some point (they always diverge eventually). This sort of co-optation does not interfere with the recipient organism's capacity to regulate, since the original regulative process and the new co-optive one are compatible. Thus the visual experience of a brightly colored and highly venomous Coral Snake really does (honestly) indicate a danger to predators, and the snake and any potential predator share an interest in ensuring that the predator avoid the snake. The visual experience of the very similar coloration of the non-venomous Scarlet King Snake may provoke avoidance behavior in predators, but the snake's effect on the indicator here is dishonest: the indicator state does not correspond to real danger in the local environment. In this case the interests of snake and predator diverge, and the co-optive process of regulation does interfere with the recipient's original capacity to regulate.

Co-optation thus tends to result in selection for changes in the recipient organisms that modify the reference, reinforcing the effect of co-optation if it is advantageous to the recipient (so that predators become more responsive to the 'warning' markings of venomous snakes, for example) or reducing the effect if it is detrimental to the recipient (so that predators become less responsive to the markings of harmless snakes, i.e. better at distinguishing genuine threats from mimics.) But the degree to which the latter kind of adjustment is possible depends on the role of the regulated variable in the reproductive success of the organism. If the regulation that the system provides is not highly consequential for the fitness of the organism it may be possible for selection to change the reference sufficiently (perhaps eliminating it altogether) to protect effectively against costly co-optation. But if the regulation is vital for reproductive success, the reference cannot be changed sufficiently to block exploitation—this is why manipulations such as brood parasitism and mimicry of dangerous organisms are so effective. Host birds presented with a gaping chick in the nest, or predators presented with a brightly banded snake, may be obliged to err on the side of caution.

The co-optation of regulation thus produces selective pressures on the recipient organisms. But the recipient organisms also exert selection pressures on the agents. Systems that regulate behavior are not only vulnerable to exploitation, but are often so structured that they actively *select for* the capacity to exploit them. In the first place, they select for successful co-optation. These systems work by making discriminations. A system that regulates the behavior that its possessor expresses in interacting with other organisms does so by discriminating among the organisms that it encounters and modifying the interactive behavior accordingly. The differences in behavior that result are likely to be consequential for the fitness of the organisms that elicit them: if the recipient organism is using the discriminations made by its indicator to regulate behaviors involved in foraging, predator avoidance, mate choice or care of offspring, the discrimination itself becomes a form of natural selection upon agents, where what is selected for is relative success in co-opting the regulating system. And when the interests of the agent and the recipient diverge, the recipient's discriminations select agents for the capacity to succeed in exploiting them. The bird that is less likely to eat a harmless insect resembling a toxic one, the host bird that is more likely to incubate a cuckoo egg resembling its own eggs, and the parent that is more likely to give more than its fair share to a crying offspring, all are selecting among the organisms they interact with for success in exploitation.

Behavioral levers

If a behavioral lever—a system that regulates behavior—that is both open to exploitation and worth exploiting is maintained in a population, it is probably serving some adaptive function for its possessors. Levers that are without significant adaptive value for their possessors *and* are readily and profitably exploitable by other organisms will tend to come to be exploited, in the long run, and so to be selected against. So if such a lever is maintained, there is probably some cause maintaining it. The cause could be that the lever has some adaptive value that outweighs or at least balances its costs, as in the case of the exploitable but vital systems that direct birds to feed gaping chicks in their nest or predators to avoid red-and-black banded snakes as well as most of the preferences and aversions that guide animals' behavior in foraging, predator-avoidance and social interaction. It could be that the lever is not itself adaptive but is developmentally or genetically linked to an adaptive feature, as the by-product account says is the case with the female orgasm (Lloyd 2005, p. 163). But the reference that determines what state the regulative system will tend to produce or maintain needs to be flexible and so is usually easily modified by selection or learning, so even if the rest of the system is maintained through linkage or developmental connection, selection will commonly be able to render the lever inoperative if its net cost is significant. If a behavioral lever that is accessible to exploitation and worth exploiting is selectively neutral for its possessor, this state of affairs is thus likely to be an achievement of selective evolution rather than an instance of a selective 'don't care.'

Because of their special role in regulation, references themselves or features that determine the values of new references through learning are the most likely to be modified in the evolution of regulative systems. These features—including behavioral references such as preferences, attractions, aversions, innate behavioral stimulus–response patterns, and the innate patterns of pleasure and pain responses that direct the learning of new references—thus pose a distinctive problem for evolutionary explanation. Though they may originate either by direct natural selection or by strictly non-adaptive means, the evolutionary explanation of their presence is unlikely to be either a strictly non-adaptive story (of linkage, developmental side-effects or inertia, or chance, for example) or a simple story of problem-solving adaptation. It is likely instead to be a story much better treated in the terms of extended adaptationism.

Behavioral levers and extended adaptation

What does extended adaptationism tell us about the evolution of features that contribute to behavioral regulation systems? Begin by returning to Lewontin's question: "*What* is adapting to *what*?" There are parts of the biological world for which this question is difficult to answer: *contested zones* in the Darwinian struggle for existence that are causally influenced (or in some cases regulated) by more than one organism in ways that have fitness consequences for those organisms—where one organism's niche construction impinges on another organism's body or its constructed niche, or where the extended phenotypes of genes in two different organisms overlap. In these zones, two (or more) organisms (or perhaps other functional systems) are engaged in a process in which *each is adapting to the other*, producing a feedback loop. In cases where one or both of the organisms are regulating the state of the contested zone, a special and important kind of feedback process can emerge in which *each organism is adapting to the other but is also actively 'adapting' the other to itself*, in a complex feedback process that can have remarkable evolutionary effects.

Consider first two populations of organisms, each having an effect on its own environment (perhaps directly on the other population) that has a selective effect on the other, but neither regulating that effect.[2] Here we have a simple feedback process such as those involved in many arms races. Take running speed, for example. An organism regulates its speed of locomotion at any one time, but its maximum running speed for any given terrain, a factor in some predator-prey interactions, is a constraint on the regulated locomotive speed that is not itself regulated. In the arms-race between pursuing cheetahs and fleeing antelopes, the antelopes are adapting to the maximum speed of the cheetahs; the cheetahs are adapting to the maximum speed of the antelopes. This is a form of feedback and does produce some of the distinctive results of feedback, including the possibility of relatively rapid and sustained change in one direction through the 'Red Queen' effect (Van Valen 1973)—producing, in this case, very fast cheetahs and very fast antelopes. But unless some other factor intervenes, the direction of change will remain relatively stable, so this sort of feedback tends to result in relatively straightforward adaptive evolution that produces extreme versions of common solutions to common problems: fast predators to catch prey, fast prey to escape predators, hard shells to protect against predators, strong jaws to break shells, tall trees to reach above the shade of neighbors, and so on.

Now consider a population of organisms that regulates behavior interacting with a population of co-opters that do not regulate the co-opting process: this produces a more complex feedback process characteristic of many plant–insect interactions, among others. In the interaction between a flowering plant and an insect pollinator, for example, the insect regulates its behavior, including which flowers it visits and what to do when visiting them, in accord with visual and olfactory references and what its indicators tell it about its surroundings, including the flowers. The plant is thus able to co-opt the insect's behavior to contribute to its own reproductive functions. In the coevolution of bee and orchid, the bee is adapting to the appearance and scent of the orchid, evolving new and more finely-tuned references, for example; the orchid is adapting to the *behavior that the bee produces in response* to the appearance of the orchid. (Many animal–animal interactions have this structure as well: those in which one animal's behavior is tuned to another's non-behavioral features, such as size, color, markings, etc.)

These cases can produce marked feedback effects, including many well-known forms of coevolution, but the feedback process is still relatively simple. It seems right to say the bee is solving the problem of how to find food (of the appropriate type), and the plant is solving the problem of attracting an insect that's solving the problem of how to find food (of that type). This sort of evolution often produces exaggerations or modifications of the features that the reference was originally tuned to (larger flowers, more brightly colored plumage, sweeter fruit, etc.), or potentially elaborate forms of mimicry or camouflage. But though the particularities change, the basic problems don't change much here—how to attract, how to repel, how to find, how to elude, how to choose (the best mate, the best food source).

Remarkable effects can be produced, however, when both systems are doing regulation in the contested zone. Consider interactions between organisms each of

---

[2] It should be borne in mind throughout that although I will write about the interactions as taking place between members of two discrete populations, usually of different species, the same kinds of processes may take place among members of a single population.

which is capable of regulating behavior, and in which (at least) one is co-opting a system of behavioral regulation in the other *in a regulated way*. Here the recipient organisms are regulating their own behavior, but the agent organisms are using their means of co-optation *to regulate that same recipient behavior*. Many animal social interactions and close behavioral interactions between species are of this type. In such an interaction, what the recipient organism does depends on what the agent organism does (i.e. what input it provides to the recipient's indicator) but the reverse is true as well, since the agent is regulating its own behavior according to the behavior of the recipient. If the behavior of each organism has fitness consequences for the other, we have a situation in which there is feedback not only between the selective effects that the organisms have on one another, but between the behaviors that mediate those effects.

Before I discuss an example of this two-tiered feedback structure, consider the following striking implications:

Suppose recipient organism R has a co-optable regulative system. By providing the right indicator input, agent organism A can:

M1: *modify* R's immediate behavior
M2: *modify* R's reference structure (and so R's *pattern* of behavior) by means of the mechanisms of learning (“conditioning”), including
  (a) changing or adding subsidiary references that determine the means employed to regulate the original variable
  (b) changing or adding new high-level references (such as those involved in affectional bonds, especially) (Bowlby 1969, 1977)
D1: *discriminate* among Rs on basis of responsiveness to direct modification of behavior OR of capacity to produce preferred behavior
D2: *discriminate* among Rs on basis of capacity to acquire
  (a) new behavior patterns
  (b) new affectional bonds

If the interaction has fitness consequences for R, we may expect Rs to experience selection pressure as a result of the various kinds of discrimination involved in A's co-optation. As we are presuming that the causal mechanisms employed by A in all these forms of co-optation are behavioral and so are regulated, some discrimination is involved in M1 and M2 which may have fitness consequences for R. In D1 and D2, the process of discrimination itself is regulated by A.

The distinctions among these different co-optive possibilities may be clarified by a look at a familiar example; a common pattern of interaction between humans and domestic dogs is especially easy to grasp because it (often) involves conscious intention on the part of the human participant. Note, however, that this aspect is not a necessary condition for any of the effects that I describe.

Dogs like the flavor of cheese; that is, the flavor of cheese is close to some reference that dogs use in selecting food. Humans can take advantage of this fact about dogs in several different ways. If I want to *get* a dog to jump up in the air, I can do so by holding a piece of cheese in the right place above the dog's head. If I want to *train* a dog to perform some particular behavior (jumping on command, say) I can do so by using cheese as a “reward”: the dog adds a new sub-reference to those it employs in regulating food intake, a new way of getting the food it prefers. If I want to get a dog to *like* me, and so ultimately (for example) to be willing to learn new commands from me quickly, willing to take risks to protect me, etc., giving it a lot of pieces of cheese (under particular conditions) is not a bad way to start. Moreover, if I want to *pick out* the dogs that are the easiest to get to jump, those that can jump highest, those that learn commands most quickly, or those that make friends most easily, I can do so with the help of a lot of bits of cheese. A familiar style of folk-account of the domestication of wild canines involves all of these various forms of co-optation, using food and other references (Schleidt 1998; see also Mitchell and Thompson 1990, 1991).

This is less than half the story, however. Consider the other side of the interaction. If any of these forms of co-optation of R provides a fitness advantage for A, then the As will be selected for the capacity to perform it systematically, i.e. in a regulated fashion. To do this requires that A possess a regulative system that is sensitive to success in obtaining the ‘right’ behavior from R, and to other indications of success in providing the ‘right’ input to R’s indicator for obtaining that behavior. But that means a regulative system that is almost certainly both co-optable and worth co-opting by R, unless it is exploitive. It is co-optable because it is tuned to R’s behavior; it is worth co-opting because it is consequential for R. If A’s signal is honest, then the Rs will be selected for responding selectively, and for making clear indications when A provides the ‘right’ input—e.g. expressions of pleasure or satisfaction. Using these means, R can now co-opt A’s co-opting behavior, in ways analogous to those sketched above for A (M1, M2, D1 and D2), creating a feedback loop. (If A’s signal is dishonest, the Rs will simply be selected for being unresponsive to A’s initial co-optation.)

So in the human–dog interaction, if the humans obtain fitness benefits from modifying the behavior of dogs, then they will be selected for being good at doing so, that is, for regulating the processes by which they do it. That means that humans will come to choose what to do in accord with what the dogs do. In particular it means that the humans will come to *like* getting dogs to perform the behavior that benefits humans, to *like* interacting with dogs in the distinctive way that co-optation demands, and to *like* succeeding in obtaining a pleased response from a dog. And *this* means that dogs can shape their own behavior (and are likely to learn or be selected so to shape it) so as to co-opt humans’ behavior and so maximize their benefit from the interaction. If a human is benefiting from being able to get a dog to jump, the dog can take advantage of this fact. It can refuse to jump except when a particularly large piece of cheese is on offer, thereby getting the human to hand over a big piece of cheese this time or training the human to provide larger pieces of cheese in general. It can jump extra frequently or extra high, encouraging the human to stay around. It can choose which humans to spend time with, according to their cheese-supplying capacity or their responsiveness to interaction. It can also benefit from humans’ learned or evolved liking for seeing a happy dog, wagging its tail to encourage more cheese-disbursement (immediate or in general) and stronger feelings of affection. And these activities, in turn, open new opportunities for co-optation by humans.

On both sides of this evolutionary interaction, references have been exapted from use in intraspecific social interactions. Both canines and humans have large repertoires of social references that are easily transferred from intraspecific to interspecific contexts: humans’ ability to distinguish emotional states in other humans was readily exapted to serve in interactions with other mammal species, for example. Liking to see a smiling family member is easily adjusted to include (a more modest level) of liking to see a smiling stranger, or a grinning tail-wagging dog.

Explaining behavioral levers

Behavioral levers—especially those involved in mutual behavioral co-optation—pose a special problem for evolutionary explanation. The feedback processes involved in unregulated interactions such as simple arms races, and those involved in the more complex interactions that are regulated on one side only, do pose some problems for adaptationist explanation, insofar as the details of the ‘problem’ solved by the levers or other features involved in the interaction change as the evolutionary process proceeds. But in interactions involving mutual behavioral co-optation—regulation in the contested zone—the processes of evolution include complex feedback loops that make it possible for organisms to act upon their environments in such a way as to bring into being wholly new *kinds* of problems. Today’s bee is using its preference for particular features of modern flowers to solve the ancient problem of how to find food. It uses many other references, though—those governing behavioral interactions—to solve new problems that evolved along with their solutions: the problem of how to communicate the location of a food source to hivemates, for example, or the problem of how to drive off a hungry bear. In such cases, there is no straightforward answer to the question ‘‘What is adapting to what?’’ Interacting dogs and humans, or interacting honeybees, are adapting to one another and co-opting one another in a tangled web of feedback relations between immediate behavioral responses and learned and evolved patterns. For these cases, the sort of simple problem-solving adaptive explanation characteristic of simple adaptationism is clearly inadequate. But on the other hand these features are quite unlikely to be maintained in a population in which they serve no adaptive purpose, because they are so structured as to tend to select for co-opters that will in turn select either for or against the original features. A behavioral lever that is employed in behavioral interaction is *apt for aptation*—‘ready to hand’ for further adaptation, for exaptation, or for co-optation. As a result, these levers—features that contribute to systems of behavioral regulation both vulnerable to co-optation and worth co-opting—constitute a category of features ill-suited to the anti-adaptationist’s assumption that anything that isn’t a simple problem-solving adaptation is non-adaptive. Although this category may be a very small one numerically, it is one of distinctive importance and includes some of the features that are most often the focus of evolutionary controversy: those that govern the intimate social relations of human beings as well as other animals, including those involved in the regulation of mating behavior, intraspecies aggression, and parent–offspring interactions.

How, then, ought we to proceed in explaining these special biological features? There are two basic forms that an explanation using feedback or circular causality might take. One is an explanation that describes the full evolutionary path traveled by the interacting populations, mapping the interacting processes of response and selection and the feedback circuits that they create. Such explanations would reduce the feedback processes of circular causality to chains of linear causal relations; they are rarely likely to be possible in practice, but the in-principle possibility of constructing them is important for the second sort of explanation. The other sort of explanation will depend on developing models of a variety of *kinds* of processes of selective interaction, according to the kinds of interaction involved, the kinds of regulation each participant is engaged in, and the consequences of these for both participants. The best place to start, I think, is by combining Lewontin’s question and the notion of regulation. *What* is adapting to *what*? What is *regulating* what? By what means? With what consequences? When we can give answers—and they will most often be non-straightforward answers—to these questions, we will have the resources to understand (though probably not trace) the evolutionary path the interacting populations followed by understanding the basic structure of the interaction.

The non-straightforward explanations of circular causality often seem unsatisfying. ‘‘I am here because of you, and you are here because of me,’’ says Mulla Nasrudin in the tale with which we began, knowing that the travelers will not feel that their question has been answered. The reader, of course, has been given the whole story of this simple interaction, and could also easily assimilate this case to a simple model of ‘flight and pursuit’ interactions. But though knowing the whole sequence of events, in this simple case as well as in much more complex evolutionary interactions, does succeed in reducing the causal circles to causal chains, it does not quite reduce non-straightforward explanations to straightforward ones. A straightforward explanation, adaptive or non-adaptive, has a beginning point in one of the basic—i.e. general—problems posed by evolutionary environments or one of the general rules about non-adaptive processes in evolution. A non-straightforward explanation, even one that specifies all the steps in the interaction, always starts *in medias res*, always leaves unanswered at some point the further question ‘‘and why was *that* the problem needing solution just then?’’, and so is always in that sense irreducibly historical.

Implications for extended adaptationism

What implications does the existence of behavioral levers and their role in the regulation of behavior have for our understanding of extended adaptation? I'll address this question by considering the three extended-adaptationist approaches discussed above: coevolution, the extended phenotype, and niche construction.

The most important implication is simply that the special causal structures that distinguish regulation from unregulated processes are too consequential to be ignored. All three approaches essentially do ignore this distinction. Coevolutionary theory draws no systematic distinction between the selective contributions of regulated and unregulated effects that organisms have on one another. In his treatment of extended phenotypes, Dawkins (1982) describes both regulated and unregulated effects that genes in one organism can have on another as ‘‘manipulation,’’ obscuring the distinction between them, though he does emphasize the distinction between extended adaptations and mere effects (1982, 2004). Niche construction theory (Odling-Smee et al. 2003) similarly fails to draw a clear distinction between regulated effects that organisms have on their environments and unregulated effects such as waste production.

The examples of ‘regulating in the contested zone’ that I have discussed are interactions between free-living animals. But the same implications apply to interactions involving plants (e.g. active defense responses and responses employed in symbiosis), fungi, and microbes, and to those between hosts and endosymbionts or endoparasites (Ham and Bent 2002; Dicke and van Poecke 2002; Soosaar et al. 2005; Ferguson and Mathesius 2003; Scheel and Wasternack 2002; Whitham and Dinesh-Kumar 2002; Barber et al. 2000; Gresshof et al. 2003; Jog and Watve 2005; Moore 2002). Moreover, if those extended adaptationists who argue for the existence of functionally-organized (i.e. self-regulating) systems above the level of the individual organism are correct, such systems are likely also to engage in this sort of regulative conflict. Functional systems at different levels within the same organism or larger system often have selective interests that intersect in interesting and consequential ways including both commonalities and conflicts (Buss 1987; Dawkins 1982; Sober and Wilson 1998); they may often also engage in mutual co-optation by means of regulation in ways that contribute significantly to the evolution of complexity, both organismal and social (Maynard Smith and Szathmary 1995; Holland 1995). Altruistic behavior, for example, may be understood as the result of a complex interplay of individual adaptation, co-optation by other individuals and co-optation by a functionally-organized group, so that the behavioral effectors of a single organism are being recruited to regulate a diverse array of variables in other organisms and in the social group. The resultant selective feedbacks (on social instincts and signaling behavior, for example) may have been instrumental in the evolution of complex sociality.

True social learning gives humans, especially, an unparalleled capacity for adjusting their behavioral references in adaptive ways by learning new skills, practices and useful beliefs (Richerson and Boyd 2005). The same capacity also opens up unique possibilities for both co-operative and exploitive co-optation, by enabling humans to modify one another's behavioral *references* by modeling or teaching.

Biological levers are the focus of selective interactions in the contested zone under regulation by both systems, and pose a special challenge for evolutionary explanation. Extended adaptationist approaches to evolutionary explanation are far better fitted to handle the task of explaining their evolution than are the 'straightforward' approaches of simple adaptationism or anti-adaptationism. The three extended adaptationist approaches canvassed above have different strengths and weaknesses. Each of these approaches, enriched by attention to the general role of regulation and the special effects of regulative interaction, has the potential to make a distinct contribution to this explanatory task.

The case of the female orgasm

Debate about how to explain the female orgasm, as Lloyd (2005) shows, played a crucial role in shaping the way that both adaptationists and their critics think about adaptation. Recognizing the role of biological leverage, and the complex selective forces that it creates, suggests the need to rethink this important and controversial case. Begin with the presumption—which no-one disputes—that aspects of the male orgasm serve as a 'reward' adapted to guide male behavior. If this is the case, those features constitute a behavioral lever. If those features, selected for in the male, appear in the female as a developmental by-product, both sexes now possess similar behavioral levers. What effect do these have on behavior? We have very little idea. (*Exactly* what behavior the male orgasm was originally adapted to control, and in what manner, is worth a great deal more consideration than it has received.) It seems likely, though, that for both sexes patterns in the distribution and quality of orgasms they experience has some effect on their mating behavior or other responses to past or potential mates—that is, that orgasm in fact makes some kinds of discrimination, and in the male is adapted to do so. But if this is the case, the lever is almost certainly both worth co-opting and worth protecting from co-optation, for many of the details of behavior in and immediately surrounding mating are highly consequential for the fitness of both partners. Moreover, the lever is likely to be highly co-optable, since it is making discriminations precisely on the basis of the experience its possessor has in interacting with a mate. It is thus eminently apt for exaptation and for both co-operative and exploitive co-optation. Moreover, if success in co-opting a mate's orgasm is of significant benefit to one or both sexes, this may produce a secondary behavioral lever in the form of a preference for mates that show behaviors characteristic of (successfully co-opted) orgasm, which in turn would open up the possibility of further exaptation or co-optation.

The aspects of orgasm that are of most interest to the many participants in the debate are those that produce behavioral leverage: pleasurable sensations and emotional effects. The explanation of how these features evolved—in both male and female cases—is therefore very likely to be a complex story about adaptation, exaptation and co-optation, even though Symons (1979) and Lloyd (2005) are likely correct that the initial step in the evolution of the female orgasm was a byproduct of the early evolution of the male orgasm. Both non-adaptive and simple adaptive stories are likely to be deeply wrong. The obvious capacity of some elements of the complex suite of features we call orgasm to contribute to behavioral regulation may explain the determination of thinkers (over hundreds of years, at least) to find that female orgasm is *for* something even though it has always been mysterious what that something might be. Levers in adaptive systems are the paradigm of things that are 'for something'—they are for regulating something—yet *what* they are for is apt to be changed, either from within by their possessors or from without by co-opting others.

References

Alcock J (1987) Ardent adaptationism. Nat Hist 96(4):4

Ashby WR (1956) An introduction to cybernetics. John Wiley & Sons, New York

Barash DP (2005) Let a thousand orgasms bloom! a review of the case of the female orgasm by Elisabeth A. Lloyd. Evol Psychol 3:347–354

Barber I, Hoare D, Krause J (2000) Effects of parasites on fish behaviour: A review and evolutionary perspective. Rev Fish Biol Fish 10:131–165

Bateson P (1988) The active role of behaviour in evolution. In: Ho M-W, Fox SW (eds) Evolutionary processes and metaphors. Wiley, New York, pp 191–207

Bernard C (1865) An introduction to the study of experimental medicine, English translation Greene HC (trans.) 1927. Macmillan, New York

Bowlby J (1969/1997) Attachment, vol 1 of Attachment and loss. Pimlico, London. Reprint of Hogarth, New York

Bowlby J (1977) The making and breaking of affectional bonds. Br J Psychiatry 130:201–210; 421–431

Boyd R, Richerson PJ (1985) Culture and the evolutionary process. University of Chicago Press, Chicago

Buss LW (1987) The evolution of individuality. Princeton University Press, New Jersey

Cavalli-Sforza LL, Feldman MW (1981) Cultural transmission and evolution: a quantitative approach. Princeton University Press, Princeton

Cziko G (2000) The things we do: using the lessons of Bernard and Darwin to understand the what, how, and why of our behavior. MIT Press, Cambridge, MA

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

Dawkins R (1986) The blind watchmaker. Norton, New York

Dawkins R (2004) Extended phenotype – but not too extended. A reply to Laland, Turner and Jablonka. Biol Philos 19:377–396

Day RL, Laland KN, Odling-Smee J (2003) Rethinking adaptation: the niche-construction perspective. Perspect Biol Med [Chicago] 46(1):80–95

Dennett DC (1995) Darwin's dangerous idea. Simon and Schuster, New York

Dicke M, van Poecke RMP (2002) Signalling in plant-insect interactions: signal transduction in direct and indirect plant defence. In: Scheel D, Wasternack C (eds) Plant signal transduction. Oxford University Press, Oxford, pp 289–316

Ehrlich PR, Raven PH (1964) Butterflies and plants: a study in coevolution. Evolution 18:586–608

Ferguson BF, Mathesius U (2003) Signaling interactions during nodule development. J Plant Growth Regul 22:47–72

Fisher J, Hinde RA (1949) The opening of milk bottles by birds. Br Birds 42:347–357

Futuyma DJ, Slatkin M (eds) (1983) Coevolution. Sinauer Associates, Sunderland, MA

Godfrey-Smith P (2001) Three kinds of adaptationism. In: Orzack SH, Sober E (eds) Adaptationism and optimality. Cambridge University Press, Cambridge, pp 335–357

Gould SJ (1987) Freudian slip. Nat Hist 96(2):14–21

Gould SJ, Lewontin RC (1979) The spandrels of San Marco and the panglossian paradigm: a critique of the adaptationist program. In: Proceedings of the Royal Society, London, vol 205, pp 581–598, reprinted in Sober E (ed) 1994, Conceptual issues in evolutionary biology, 2nd edn. MIT Press, Cambridge MA, pp 73–90

Gould SJ, Vrba E (1982) Exaptation – a missing term in the science of form. Paleobiology 8:4–15

Gresshoff PM, Rose RJ, Singh M, Rolfe BG (2003) Symbiosis signals. Today Life Sci May/June, 30–33

Griffiths PE, Gray RD (1994) Developmental systems and evolutionary explanation. J Philos 91(6):277–304

Ham JH, Bent A (2002) Recognition and defense signalling in plant/bacterial and fungal interactions. In: Scheel D, Wasternack C (eds) Plant signal transduction. Oxford University Press, Oxford, pp 198–225

Holland J (1995) Hidden order: how adaptation builds complexity. Addison-Wesley, Reading, MA

Jablonka E (2001) The systems of inheritance. In: Oyama S, Griffiths PE, Gray RD (eds) Cycles of contingency. MIT Press, Cambridge, MA, pp 99–116

Jablonka E, Lamb MJ (1995) Epigenetic inheritance and evolution: the Lamarckian dimension. Oxford University Press, Oxford

Janzen DH (1980) When is it coevolution? Evolution 34:611–612

Jog M, Watve M (2005) Role of parasites and commensals in shaping host behaviour. Curr Sci 89:1184–1191

Keller EF (2002) Making sense of life: explaining biological development with models, metaphors, and machines. Harvard University Press, Cambridge, MA

Laland KN (2004) Extending the extended phenotype. Biol Philos 19(3):313–325

Laland KN, Odling-Smee FJ, Feldman MW (1996) On the evolutionary consequences of niche construction. J Evol Biol 9:293–316

Laland KN, Odling-Smee FJ, Feldman MW (2000) Niche construction, biological evolution, and cultural change. Behav Brain Sci 23(1):131–175

Lewontin RC (1983a) Gene, organism, and environment. In: Bendall DS (ed) Evolution from molecules to men. Cambridge University Press, Cambridge, pp 273–285

Lewontin RC (1983b) The organism as the subject and object of evolution. Scientia 118:65–82

Lewontin RC (1984) Adaptation. In: Sober E (ed) Conceptual issues in theoretical biology: an anthology. MIT Press, Cambridge, MA, pp 234–251

Lloyd EA (2005) The case of the female orgasm: bias in the science of evolution. Harvard University Press, Cambridge, MA

Maynard Smith J, Harper D (1995) Animal signals: models and terminology. J Theor Biol 177:305–311

Maynard Smith J, Szathmary E (1995) The major transitions in evolution. W.H. Freeman, Oxford

Mayr E (1983) How to carry out the adaptationist program? American Naturalist 121:324–333, reprinted in Mayr E, Towards a new philosophy of biology. Harvard University Press, Cambridge, MA, 1988, pp 149–159

Mitchell RW, Thompson NS (1990) The effects of familiarity on dog-human play. Anthrozoos 4(1):24–43

Mitchell RW, Thompson NS (1991) Projects, routines, and enticements in dog-human play. In: Bateson PPG, Klopfer PH (eds) Perspectives in ethology, vol 9. Plenum Press, New York, pp 189–215

Moore, J (2002) Parasites and the behavior of animals. Oxford University Press, Oxford

Odling-Smee FJ, Laland KN, Feldman MW (2003) Niche construction: the neglected process in evolution. Princeton University Press, Princeton, NJ

Oyama S (2000) The ontogeny of information: developmental systems and evolution, 2nd edn. Duke University Press, Durham, NC

Oyama S, Griffiths PE, Gray RD (eds) (2001) Cycles of contingency. MIT Press, Cambridge, MA

Plotkin HC (ed) (1988) The role of behavior in evolution. MIT Press, Cambridge, MA

Preston, B (1988) Why is a wing like a spoon? A pluralist theory of function. J Philos 95(5):215–254

Richerson PJ, Boyd R (2005) Not by genes alone: how culture transformed human evolution. University of Chicago Press, Chicago

Robertson DS (1991) Feedback theory and Darwinian evolution. J Theor Biol 152: 469–484

Scheel D, Wasternack C (2002) Signal transduction in plants: cross talk with the environment. In: Scheel D, Wasternack C (eds) Plant signal transduction. Oxford University Press, Oxford, pp 1–5

Schleidt WM (1998) Is humaneness canine?. Hum Ethol Bull 13:1–4

Shah I (1972) The exploits of the incomparable Mulla Nasrudin. E.P. Dutton, New York

Sober E, Wilson DS (1998) Unto others: the evolution and psychology of unselfish behavior. Harvard University Press, Cambridge, MA

Soosaar JLM, Burch-Smith TM, Dinesh-Kumar SP (2005) Mechanisms of plant resistance to viruses. Nat Rev Microbiol 3:789–798

Symons D (1979) The evolution of human sexuality. Oxford University Press, New York

Thomas F, Adamo S, Moore J (2005) Parasitic manipulation: where are we, and where should we go? Behav Processes 68:185–199

Thompson JN (1995) The coevolutionary process. University of Chicago Press, Chicago

Turner JS (2000) The extended organism: the physiology of animal-built structures. Harvard University Press, Cambridge, MA

Van Valen L (1973) A new evolutionary law. Evol Theory 1:1–30

Whitham SA, Dinesh-Kumar SP (2002) Signaling, in plant-virus interactions. In: Scheel D, Wasternack C (eds) Plant signal transduction. Oxford University Press, Oxford, pp 226–249

Wiener N (1961) Cybernetics or control and communication in the animal and the machine, 2nd edn. MIT Press, Cambridge, MA

Wolf JB, Brodie ED III, Cheverud JM, Moore AJ, Wade MJ (1998) Evolutionary consequences of indirect genetic effects. Trends Ecol Evol 13(2):64–69