# Niche Construction and Cognitive Evolution

**Benjamin Kerr**
Department of Biology
University of Washington
Seattle, WA, USA
kerrb@u.washington.edu

## Abstract

Despite the fact that animal behavior involves a particularly powerful form of niche construction, few researchers have considered how the environmental impact of behavior may feed back to influence the evolution of the cognitive underpinnings of behavior. I explore a model that explicitly incorporates niche construction while tracking cognitive evolution. Agents and their stimuli are modeled as coevolving populations. The agents are born with "weights" attached to behaviors in a repertoire. Further, these agents are able to change these weights based on previous success and an inherited learning parameter. Both the agent and the stimulus receive payoffs through a behavioral interaction (where the payoff structure is influenced by the "genotype" of the stimulus). The behaving agent exhibits niche construction through its effects on stimuli (the "environment"), which can feed back to influence the value of different cognitive strategies. Here I focus on two forms of niche construction: (1) the stimulus and responding agent have common interests (positive niche construction) and (2) the stimulus and agent have dissimilar interests (negative niche construction). The form of niche construction qualitatively affects cognitive evolution (i.e., the initial behavioral probability distribution and the value of the learning parameter). Given a mutualism between the stimulus and responding agent, rapid learning and "fixed" behavioral distributions (i.e., most of the weight on a single behavior) evolve. Given an antagonism between the stimulus and agent, slower learning and "flexible" behavioral distributions (i.e., equal weight on different behaviors) evolve. I discuss these results in light of findings from the fields of ethology, psychology, and evolutionary ecology.

## Keywords

antagonism, Baldwin effect, coevolution, feedback, Goldilocks principle, innate behavior, Law of Effect, learning, mutualism, niche construction

July 8, 2007; accepted October 5, 2007

An unavoidable consequence of being alive is an alteration of one's surroundings. All organisms take up critical resources, generate detritus, and affect one another via ecological interactions. Through their physiology, development, and behavior, organisms alter the world in which they live, a process labeled "niche construction" (Lewontin 1978, 1982, 1983; Laland et al. 1996, 1999; Odling-Smee et al. 2003; Laland and Sterelny 2006). The central idea is that niches are not simply "out there" to be filled, but rather are forged in part by the activities of organisms—the "niche constructors." As organisms change their abiotic and biotic environments, they are simultaneously changing current and future selective pressures. In this way, niche construction feeds back to affect the evolution of the constructor. Such evolutionary feedback has been shown to have strong effects in a variety of biological systems, including toxin-producing bacteria in microbial biofilms (Kerr et al. 2002; Kerr 2007), flammable plants in fire-prone ecosystems (Kerr et al. 1999; Schwilk and Kerr 2002; Schwilk 2003), mound-building termites (Turner 2000; Odling-Smee et al. 2003), and pathogens in human populations using antibiotics (Boni and Feldman 2005). In this article, I explore the effects of niche construction on cognitive evolution in animals.

Animal behavior is an especially potent form of niche construction. Foraging, mating, artifact construction, and many other behaviors involve movement within the environment, response to the environment, and active manipulation of the environment. Animals choose the stimuli to which they respond and potentially alter their stimulus pool through their responses. For instance, when a gray lag goose exhibits stereotyped egg-rolling behavior (the nest-bound goose extends her neck and gently rolls a stray egg back into the nest; Lorenz and Tinbergen [1938] 1970), she places the stimulus (her egg/offspring) in a position of protection (the nest). As another example, when an osprey learns to catch fish (Edwards 1989; Mameli and Bateson 2006), these food items are removed from the stimulus pool. Such alterations set up feedback cycles that could theoretically affect the evolution of the cognitive structure supporting different behavioral responses. In the examples above, one behavior (egg-rolling) helps the stimulus (egg), whereas another behavior (fish-catching) hurts the stimulus (fish). Does the nature of the interaction between the stimulus and the responding organism influence the nature of the response?

Interestingly, much of the theoretical work on animal behavior has not incorporated niche construction. Rather, the approach has been largely "externalist" (*sensu* Godfrey-Smith 1998) in that the cognitive strategies of behaving organisms are explained by referring to properties of an autonomous external world. One example of this externalism comes from the theoretical literature on animal learning. The upshot of many mathematical models of learning has been called the "Goldilocks principle" (Kerr and Feldman 2003). If the environment changes too quickly, then what was learned yesterday has no use today and learning is simply a waste of time. On the other hand, if the environment changes too slowly, the organism should be "born ready" with the correct response and avoid the costly learning process altogether. The conclusion is that environmental variability needs to be "just right" at some intermediate level to favor the learning process. These theoretical models are externalist because the explanation of the cognitive strategy emanates from the properties of its external environment. It is the changeability of an autonomous world that dictates the course of cognitive evolution (Arnold 1978; Plotkin and Odling-Smee 1979; Johnston 1982; Stephens 1987, 1991; Bergman and Feldman 1995; Feldman et al. 1996; Godfrey-Smith 1998; Kerr and Feldman 2003).

In his book, *Thought in a Hostile World*, Kim Sterelny offers a different perspective. Sterelny explicitly considers how the effects organisms have on their environments feed back to influence their cognitive strategies. This requires the recognition that a behaving agent does not only affect its own fitness through its behavior, but also affects the welfare of its stimulus. That is to say, behavior can involve a *biological interaction*, simultaneously affecting the environmental stimulus and the responding organism and setting the stage for coevolutionary feedback: "Biological interactions change the informational character of environments, and these changes often select for more complex tracking and control systems. . . agents' tracking capacities have coevolved with each other and their environments" (Sterelny 2003: 17)

Such considerations are particularly relevant when the stimulus is an organism itself (e.g., a prey item, a mate, an offspring, a competitor, a predator, etc.). Sterelny discusses how the nature of these biological interactions may influence cognitive evolution. When the behaving organism and its stimulus have an antagonistic relationship (e.g., prey–predator, host–pathogen, etc.), then behaviors in response to a particular stimulus that reward the agent select against the stimulus (e.g., a predatory success leads to a *dead* prey item). This antagonism means that the stimulus pool is continually being selected to deceive the agent. Sterelny (2003: 25) posits that these "hostile" relationships lead to lower reliability of stimuli: "Thus hostility changes the informational character of local environments, degrading the covariation between easily discriminated cues and the functional properties they signal."

In this article, I build on Sterelny's analysis to explore how the nature of the interaction between behaving organisms and their stimuli influences the evolution of cognitive strategies. The modeling framework is broad enough to capture at least some of the diversity of animal behavior. For instance, the agents in the model can be born with relatively fixed behavioral responses that are little affected by experience. Alternatively, organisms can exhibit very flexible behavioral strategies, where learning can guide the behavioral response to stimuli. By explicitly including the stimuli as a coevolving population, the model starts to address how the niche construction inherent in behavior can feed back to influence the cognitive strategy underlying the behavior itself.

## The Model

### Overview of the Model

In order to explore the evolutionary consequences of the feedback from niche construction, I construct an agent-based simulation tracking both a stimulus population ($N_S$ stimuli) and a population of organisms responding to stimuli ($N_R$ individuals)—hereafter labeled "responders." The simulations monitor stimuli and responders within and across generations. Within generations, each responder reacts to random stimuli over a series of discrete time points. For simplicity, I assume that all stimuli are perceived as equivalent by any responder (a future paper will expand this model to consider cases where different stimuli can be perceived differently). A responder can use any one of $n$ behaviors (the size of its behavioral repertoire) when reacting to a stimulus. Responder action determines the payoff to the responder and stimulus simultaneously. These payoffs accumulate throughout the lifetime and determine fitness. At the end of the lifetime ($T_S$ and $T_R$ time steps for the stimuli and responders, respectively), offspring are asexually produced in proportion to the fitness of each parental individual. Both responder and stimulus offspring can be mutants, which allow critical parameters to change over time. This simulation allows us to explore the coevolution of responders and their stimuli, specifically how the effects responders have on their stimuli feed back to influence the evolution of the responders' cognitive strategies.

### Representation of Responders

There is a total of $N_R$ responders in the simulation. At any time point $t$ within its lifetime, each responder carries a behavioral distribution, which can be represented by the vector

$$\mathbf{p}(t) = \langle p_1(t), p_2(t), p_3(t), \ldots, p_n(t)\rangle, \tag{1}$$

where $p_i(t)$ is the probability of using behavior $i$ in response to a stimulus encountered at time $t$. Note that $\mathbf{p}(t)$ is a proper probability distribution, such that

$$\sum_{i=1}^{n} p_i(t) = 1. \tag{2}$$

The process of *learning* is modeled as a change to a responder's behavioral distribution over time [i.e., $\mathbf{p}(t+1)$ need not equal $\mathbf{p}(t)$]. After each interaction with a stimulus, the behavioral distribution is altered in a way to place more weight on highly rewarding behaviors. In this way, the learning process echoes the basic sentiment of Thorndike's famous *Law of Effect*:

> Of several responses made to the same situation, those which are accompanied or closely followed by satisfaction to the animal will, other things being equal, be more firmly connected with the situation, so that when it recurs, they will be more likely to recur (Thorndike [1911] 2000)

This law was formulated after a series of experiments in which Thorndike placed various animals in uncomfortable positions (e.g., a cat inside a "puzzle box" with various levers, pulleys, and knobs). The animals gradually discovered a way to free themselves from the discomfort (e.g., the cat discovers the particular switch that opens the box). Over time, this particular escape behavior became more and more likely to occur, a process that Thorndike referred to as the "stamping in" of the behavior. Thorndike held that it was rewarding behaviors generally that were "stamped in."

Imagine that a responder at time point $t$ has employed behavior $i$ in response to a stimulus (this behavior was chosen in accordance with its behavioral probability distribution). Now, assume that the payoff of this behavior is $\pi_i(t)$. The Law of Effect is captured in the following way:

$$\mathbf{p}(t+1) = \mathbf{p}(t)\mathbf{U}(\pi_i(t)), \tag{3}$$

where $\mathbf{U}$ is the $n \times n$ "update" matrix

$$\mathbf{U}(\pi_i(t)) = \begin{bmatrix} 1-\lambda\pi_i(t) & 0 & 0 & \cdots & \lambda\pi_i(t) & \cdots & 0 \\ 0 & 1-\lambda\pi_i(t) & 0 & & \lambda\pi_i(t) & & 0 \\ 0 & 0 & 1-\lambda\pi_i(t) & & \lambda\pi_i(t) & & 0 \\ \vdots & & & \ddots & \vdots & & \vdots \\ 0 & 0 & 0 & \cdots & 1 & \cdots & 0 \\ \vdots & & & & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda\pi_i(t) & \cdots & 1-\lambda\pi_i(t) \end{bmatrix}. \tag{4}$$

Note that $\mathbf{U}$ is a stochastic matrix (each row sums to unity) and it is the $i$th column of $\mathbf{U}$ that has no zero entries. This way of representing learning is similar to the "linear operator" technique introduced by early mathematical psychologists (e.g., Bush and Mosteller 1955). The same process can be represented by focusing on the probability of any behavior $j$ (with $j \in \{1, 2, 3, \ldots, n\}$) after stimulus interaction:

$$p_j(t+1) = \begin{cases} (1 - \lambda\pi_i(t))p_j(t) & \text{if } i \neq j \\ p_j(t) + \sum_{k \neq i} \lambda\pi_i(t)p_k(t) & \text{if } i = j \end{cases}. \quad (5)$$

Thus, given that behavior $i$ is used at time $t$ and a payoff of $\pi_i(t)$ is received by the responder, the probability of the same behavior at the next point in time ($t + 1$) increases in proportion to the size of the reward. Furthermore, this extra weight is gathered from the probabilities of each of the other behaviors. The more rewarding behavior $i$, the more likely it is to occur in the future. The parameter $\lambda$ gives the responder's "speed of learning," that is, $\lambda$ measures how dramatically the behavioral distribution is reweighted after an interaction with a stimulus.

Up to this point, we have described the learning process within the lifetime of a responder (the $T_{\mathrm{R}}$ time steps over which the responder lives). In order to discuss cognitive evolution across generations, we must specify how reproduction occurs and the quantities inherited by any responder at its birth. For simplicity, I assume asexual semelparous reproduction (each offspring has a single parent from the previous nonoverlapping generation). Offspring are produced in the very last time step of parental generation (time step $T_{\mathrm{R}}$). There will be $N_{\mathrm{R}}$ responders at this time, collectively producing $N_{\mathrm{R}}$ offspring (i.e., responder population size remains constant). Any one of these offspring has the $r$th responder (with $r \in \{1, 2, 3, \ldots, N_{\mathrm{R}}\}$) as its parent with probability $\Pi_r/\bar{\Pi}$. The quantity $\Pi_r$ refers to the cumulative payoffs to responder $r$ over its lifetime

$$\Pi_r = \left(\sum_{t=0}^{T_{\mathrm{R}}-1} \pi(t)\right) e^{-c\lambda_r}. \quad (6)$$

For convenience, the subscript on $\pi$ is dropped in equation [6]. The lifetime cost of learning is represented by the factor $e^{-c\lambda_r}$, such that the cumulative payoff decreases as the cost parameter $c$ increases or as the learning parameter for responder $r$ ($\lambda_r$) increases. The quantity $\bar{\Pi}$ is simply the population average of cumulative payoffs

$$\bar{\Pi} = \frac{\sum_{r=1}^{N_{\mathrm{R}}} \Pi_r}{N_{\mathrm{R}}}. \quad (7)$$

The offspring inherits two components from its parent: (i) a starting behavioral distribution $\mathbf{p}(0)$ and (ii) a "speed of learning" parameter $\lambda$. Note that if a parent teaches its offspring what it has learned, then a form of Lamarkian inheritance applies with $\mathbf{p}_{\text{offspring}}(0) = \mathbf{p}_{\text{parent}}(T_{\mathrm{R}})$. While such cultural inheritance is extremely interesting, here I focus on the case where the initial behavioral distribution of an offspring resembles the initial behavioral distribution of its parent [i.e., $\mathbf{p}_{\text{offspring}}(0)$ resembles $\mathbf{p}_{\text{parent}}(0)$]. In this case, the genotype of responder $r$ ($\mathrm{G}_r$) is the set of its initial behavior distribution and speed of learning parameter: $\mathrm{G}_r \equiv \{\mathbf{p}_r(0), \lambda_r\}$. This genotype also defines the cognitive strategy of the responder.

The inheritance process does not involve perfect fidelity; that is, mutations in inherited quantities can occur. Imagine the parent has genotype $\mathrm{G}_j = \{\mathbf{p}_j(0), \lambda_j\}$, with

$$\mathbf{p}_j(0) = \langle p_{j,1}(0), p_{j,2}(0), p_{j,3}(0), \ldots, p_{j,n}(0)\rangle, \quad (8)$$

where $p_{j,i}$ gives the probability that individual $j$ uses behavior $i$ in response to the stimulus. The offspring's genotype $G_k = \{\mathbf{p}_k(0), \lambda_k\}$. The offspring's speed of learning is

$$\lambda_k = \begin{cases} \min(\max(\lambda_j + \delta_\lambda x, \lambda_{\min}), \lambda_{\max}) \\ \qquad\qquad \text{with probability } \mu_\lambda \\ \lambda_j \\ \qquad\qquad \text{with probability } 1 - \mu_\lambda \end{cases}, \quad (9)$$

where $\mu_\lambda$ is the probability that a mutation in the $\lambda$ parameter occurs, $\delta_\lambda$ is the maximum mutational difference between a parent and an offspring in the $\lambda$ parameter, $\lambda_{\min}$ and $\lambda_{\max}$ are the minimum and maximum $\lambda$ parameters possible for any individual, respectively, and $x$ is a random variable ($x \sim$ Unif $(-1, 1)$). In order to find the offspring's behavior distribution, consider the following weights:

$$q_{k,i}(0) = \begin{cases} \min(\max(p_{j,i}(0) + \delta_p x, p_{\min}), p_{\max}) \\ \qquad\qquad \text{with probability } \mu_p \\ p_{j,i}(0) \\ \qquad\qquad \text{with probability } 1 - \mu_p \end{cases}, \quad (10)$$

where the parameters $\mu_p$, $\delta_p$, $p_{\min}$ and $p_{\max}$ have meanings similar to those described above for $\mu_\lambda$, $\delta_\lambda$, $\lambda_{\min}$ and $\lambda_{\max}$ respectively. The offspring behavioral distribution is given by

$$p_{k,i}(0) = \frac{q_{k,i}(0)}{\sum_{i=1}^{n} q_{k,i}(0)}. \quad (11)$$

To summarize the mutational process, when the mutational probabilities ($\mu_\lambda$ and $\mu_p$) are small, the offspring is likely to resemble its parent exactly in its speed of learning and initial behavior distribution. As these mutation probabilities increase, it becomes more likely that the offspring will deviate in these quantities from its parent (although with small $\delta$ values, there will still be a tight parent–offspring correlation).

This representation of responders is broad enough to capture several phenomena of interest as subcases. For instance, consider a case where responder $r$ is born with an uneven behavior distribution in which one behavior ($i$) initially carries nearly all the weight (e.g., $p_{r,i}(0) \approx 1$ and $p_{r,j}(0) \approx 0$ for all $j \neq i$) and learning is very slow ($\lambda_r \approx 0$). This organism would exhibit one form of innate behavior (a stereo-typed response to a stimulus that is present from birth and not altered by experience). On the opposite side of the spectrum, consider a case where responder $r$ is born with an even behavioral distribution (e.g., $p_{r,i}(0) \approx 1/n$ for all $i$) and learning is faster ($\lambda_r > 0$). This organism would exhibit a flexible learning approach (response to stimuli would change through time in response to payoffs received).

### Representation of Stimuli

There is a total of $N_S$ stimuli in the simulation. While each responder can control the way that it behaves, we assume that each stimulus controls the payoff of each responder behavior, both to the responder and to the stimulus. If responder $r$ uses behavior $i$ in response to stimulus $s$ at time point $t$, the payoffs to the responder and stimulus are denoted by $\pi_{i,s}(t)$ and $\phi_{i,s}(t)$, respectively. We assume that there is some functional relationship between responder and stimulus payoff:

$$\phi = f(\pi). \tag{12}$$

For instance, if $f$ is a monotonic decreasing function ($\partial f/\partial \pi < 0$), then there is an antagonistic relationship between responders and stimuli (behaviors that are highly rewarding to the responder are not rewarding to the stimulus and vice versa). If $f$ is a monotonic increasing function ($\partial f/\partial \pi > 0$), then there is a mutualistic relationship between responders and stimuli (behaviors that are highly rewarding to the responder are highly rewarding to the stimulus and behaviors that are less rewarding to the responder are also less rewarding to the stimulus).

The stimulus does not control the shape of $f$ (although see the Discussion); however, the stimulus does define the value of responder behaviors (and thus how these same behaviors pay off to the stimulus itself through the function $f$). To make this map explicit, let $\pi_{\min}$ and $\pi_{\max}$ denote the minimum and maximum payoff to a responder respectively. The following set is $n$ evenly spaced values within the interval $[\pi_{\min}, \pi_{\max}]$:

$$\mathbf{P} = \left\{ \pi_{\min}, \pi_{\min} + \frac{\pi_{\max} - \pi_{\min}}{n-1}, \ \pi_{\min} + \frac{2(\pi_{\max} - \pi_{\min})}{n-1}, \ldots, \right. \\ \left. \pi_{\min} + \frac{(n-2)(\pi_{\max} - \pi_{\min})}{n-1}, \ \pi_{\max} \right\}. \tag{13}$$

Let the set of behaviors be

$$\mathbf{B} = \{1, 2, 3, \ldots, n-1, n\}. \tag{14}$$

The stimulus controls the one-to-one map:

$$\Gamma\colon \mathbf{B} \mapsto \mathbf{P}. \tag{15}$$

Note that there are $n!$ possible $\Gamma$ functions (or "evaluation functions"). Stimulus $s$ is born with one of these evaluation functions ($\Gamma_s$) such that responder $r$ using behavior $i$ in response to this stimulus receives the payoff

$$\pi_{i,s}(t) = \Gamma_s(i), \tag{16}$$

and the stimulus receives the payoff

$$\phi_{i,s}(t) = f \circ \Gamma_s(i), \tag{17}$$

where "$\circ$" represents functional composition. Because the evaluation function $\Gamma_s$ "belongs" to stimulus $s$, the stimulus controls the payoff structure corresponding to the behaviors used by the responder.

Asexual reproduction is assumed for stimuli. Offspring are produced in the very last time step of the parental generation (time step $T_S$). There will be $N_S$ stimuli at this time, collectively producing $N_S$ offspring. Any one of these offspring has the stimulus $s$ (with $s \in \{1, 2, 3, \ldots, N_S\}$) as its parent with probability $\Phi_s/\bar{\Phi}$. The quantity $\Phi_s$ refers to the cumulative payoffs to stimulus $s$ over its lifetime

$$\Phi_s = \sum_{t=0}^{T_S - 1} \phi(t). \tag{18}$$

For convenience, we have dropped the subscripts on $\phi$ in equation [18]. The quantity $\bar{\Phi}$ is simply the population average of cumulative payoffs

$$\bar{\Phi} = \frac{\sum_{s=1}^{N_S} \Phi_s}{N_S}. \tag{19}$$

The offspring inherits its evaluation function from its parent ($\Gamma$). As before, the inherited function is susceptible to mutation. Imagine the parent has genotype $\Gamma_j$ (the stimulus genotype is completely specified by its $\Gamma$ function). The offspring's genotype is $\Gamma_k$, where

$$\Gamma_k = \begin{cases} \Gamma_j \circ \Psi_y & \text{with probability } \mu_\Gamma \\ \Gamma_j & \text{with probability } 1 - \mu_\Gamma \end{cases}, \tag{20}$$

where $\mu_\Gamma$ is the probability that a mutation in the $\Gamma$ function occurs, and $\Psi$ is a permutation function on the set of behaviors ($\Psi : \mathbf{B} \mapsto \mathbf{B}$). In principle, $\Psi$ could be any one of the $n!$ permutation functions; however, here we focus on the subset of such functions that involve exchanging two behaviors in $\mathbf{B}$. There are $\binom{n}{2}$ such functions, which can be arbitrarily labeled $\Psi_1$, $\Psi_2, \Psi_3, \ldots, \Psi_{\binom{n}{2}}$. Finally, $y$ is a discrete random variable with $y \sim \mathrm{Unif}(1, \binom{n}{2})$. To summarize stimulus mutation, when the mutational probability ($\mu_\Gamma$) is small, the offspring is likely to resemble its parent exactly in its evaluation function ($\Gamma$) of behaviors to payoffs. As this mutation probability increases, the offspring will be more likely to deviate slightly in its map.

Overall, a responder can alter the behavior it uses in response to a stimulus (through experience during its lifetime and through mutations in learning parameters over generations) while the stimulus can alter the value of the behavior (through mutations in its evaluation function over generations). Thus, this model captures a true coevolutionary relationship between the responder and stimulus. From the responder's point of view it inherits not only its cognitive strategy from its parent, but also an affected stimulus pool. That is, the niche construction of the parental responder population leaves an ecological signature (in terms of an affected stimulus pool) that can impact the responder offspring. The effect of this "ecological inheritance" on the evolution of the cognitive strategies is the focus of this model.

## Basic Question

The central aim is to explore the effects of niche construction on cognitive evolution. Using the simulation model described above, I focus on how the nature of the relationship between the stimulus and responder (the antagonism or mutualism given by the $f$ function above) affects the evolved cognitive strategy.

## Results

Using the parameters in Table 1, I consider the evolution of the initial behavioral probability distribution of the responder [i.e., $\mathbf{p}(0)$, the distribution with which the responder is born]. It turns out that the nature of the relationship between the stimulus and responder (the form of the $f$ function) has qualitative

**Table 1.** Model parameters.

| Parameter, Variable or Function | Description | Value or Range of Values |
|---|---|---|
| $N_\mathrm{R}$ | Number of responders | 500 |
| $N_\mathrm{S}$ | Number of stimuli | 500 |
| $T_\mathrm{R}$ | Number of time steps in the lifetime of a responder | 100 |
| $T_\mathrm{S}$ | Number of time steps in the lifetime of a stimulus | 5–5000 |
| $n$ | Number of behavioral responses in the repertoire of a responder | 5 |
| $p_i(t)$ | Probability a responder uses behavior $i$ at time $t$ (can change through learning) | 0.001–0.999 |
| $\mu_p$ | Probability of a mutation in the $p_i(t)$ variables of responders during birth | 0.1 |
| $\delta_p$ | Mutational displacement "distance" in the $p_i(t)$ variables | 0.005 |
| $\lambda$ | The "speed of learning" of a responder (how quickly $p_i(t)$'s are reweighted) | 0.001–0.999 |
| $c$ | Parameter measuring how costly learning is over the responder lifetime | 0.3 |
| $\mu_\lambda$ | Probability of a mutation in the $\lambda$ variable of responders during birth | 0.0–0.1 |
| $\delta_\lambda$ | Mutational displacement "distance" in the $\lambda$ variable | 0.005 |
| $\Gamma$ | "Evaluation function" of a stimulus (which maps responder behaviors to payoffs) | $n!$ possible |
| $\mu_\Gamma$ | Probability of a mutation in the $\Gamma$ function of stimuli during birth | 0.1 |
| $\Psi$ | A permutation function that flips two elements in an ordered set of $n$ elements | $\binom{n}{2}$ possible |
| $\pi_i(t)$ | Payoff to a responder for using behavior $i$ at time $t$ (this depends on its stimulus' $\Gamma$) | 0–1 |
| $\Pi$ | Cumulative payoff to a responder over its lifetime | 0–$T_\mathrm{R}$ |
| $\phi_i(t)$ | Payoff to a stimulus when responder uses behavior $i$ at time $t$ (this depends on stimulus' $\Gamma$ and $f$) | 0–1 |
| $\Phi$ | Cumulative payoff to a stimulus over its lifetime | 0–$T_\mathrm{S}$ |
| $f$ | Global function giving the relationship between responder payoff and stimulus payoff | $\phi = f(\pi) = 1 - \pi$<br>or $\phi = f(\pi) = \pi$ |

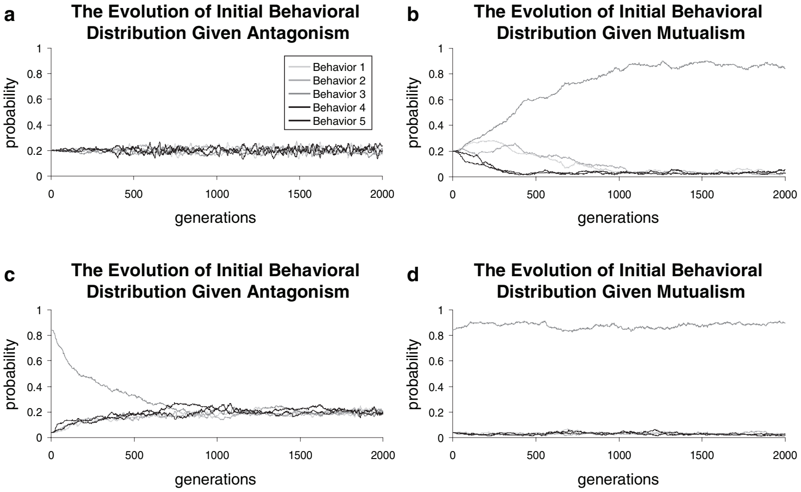

**Figure 1.**
Example runs of responder evolution. Over 2000 generations (of 100 time steps each), the evolution of the initial behavioral distribution of the responder population is tracked. For all graphs we show the average weight attached to each behavior in the responder population. Model parameters are as listed in Table 1 (with $T_S = 100$ and $\mu_\lambda = 0.1$). (a) Evolution given antagonism between stimuli and responders (here the payoff to stimulus decreases linearly with the payoff to responder: $\phi = f(\pi) = 1 - \pi$). The population of responders is initiated with a uniform initial behavioral probability distribution in the first generation. (b) Evolution given mutualism between stimuli and responders (here the payoff to stimulus increases linearly with the payoff to responder: $\phi = f(\pi) = \pi$), starting again with a uniform initial behavioral probability distribution. (c) Evolution given antagonism as in part a, however, we initialize with an uneven distribution (that is similar to the evolved distribution found in generation 2000 of part b). (d) Evolution given mutualism as in part b, however, we initialize with the same uneven distribution used in part c.

effects on the course of evolution of initial behavioral tendencies. Figures 1a and 1c track the evolution of the responder's initial distribution given an antagonistic relationship between the stimulus and responder. Specifically, $f$ is decreasing and thus behaviors highly rewarding to the stimulus are not highly rewarding to the responder and vice versa. Here, the initial behavioral distribution evolves into a near-uniform distribution. That is, all five behaviors in the repertoire of our responder are *equiprobable responses* to the stimulus at the time of the responder's birth (see Figure 2a). Figures 1b and 1d show the evolution of the responder's initial distribution given a mutualistic stimulus–responder relationship. Here, $f$ is increasing and thus the most rewarding behaviors to the stimulus are also most rewarding to the responder. In this case, one behavior evolves to possess nearly all the weight in the distribution at birth. That is, the responder is born with a *single highly probable response* to the stimulus (see Figure 2b).

Why do antagonistic stimulus–responder relationships favor the evolution of flat distributions at birth (favoring future adjustments through learning) and mutualistic relationships lead to skewed distributions at birth (favoring relatively rigid behavior)? In the case of antagonistic relations, if any responder behavior is more probable at its birth, there is selection for stimuli with an evaluation function ($\Gamma$) that maps this more probable behavior to the highest payoff for the stimulus. However, due to the antagonistic relation, what is good for the stimulus is bad for the responder (i.e., the highest stimulus payoff is simultaneously the lowest responder payoff). Thus, as the selected stimulus increases in frequency, the responder with an in-born behavioral preference is selected against. The result of this antagonistic coevolution is that the responder does not evolve behavioral preferences at birth (Figure 2a). Different starting conditions in the simulation produce the same evolutionary result in this case (compare Figure 1a to 1c). In the case of the mutualistic relation, whichever behavior is best for both stimulus and responder is weighted heavier and heavier, favoring both the rewarding stimulus and the responder with the skewed distribution simultaneously. This positive feedback leads to a situation where the responder is born with a behavioral distribution extremely skewed toward a single behavior (Figure 2b). Again, different starting conditions of the simulation lead to the same basic result (compare Figure 1b to 1d).

At this point, it is helpful to consider the diversity of the stimulus pool under different stimulus–responder relationships. Stimulus diversity is measured using the Shannon index:

$$H = -\sum_{i=1}^{n!} \gamma_i \ln \gamma_i , \qquad (21)$$

where $\gamma_i$ is the frequency of evaluation function $\Gamma_i$ in the stimulus population. Thus, this is a measure of genetic diversity of the stimuli in the context of this simulation (as $\Gamma$ is the genotype of a stimulus). Figure 3a shows that the stimulus pool is less diverse under mutualism than antagonism. This difference reflects the different feedbacks of the responder's niche construction in each case. Specifically, the responders

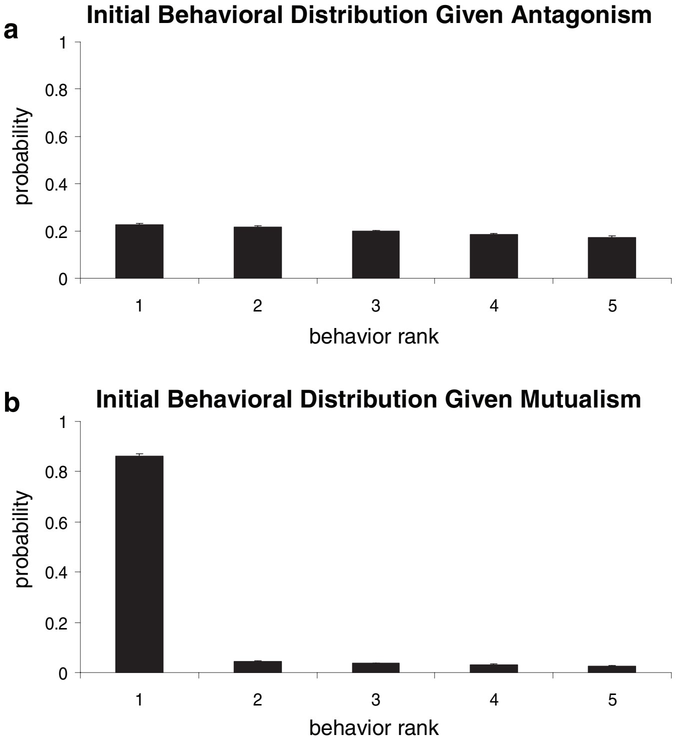

**Figure 2.**
The outcome of responder evolution. The average initial behavioral probability distribution of each of 10 separate simulation runs is recorded at generation 2500. For convenience, the behaviors are always reordered from the behavior with the most weight (behavior rank 1) to the behavior with the smallest weight (behavior rank 5). The average across 10 runs is shown (and error bars give the standard error of the mean). Model parameters are as listed in Table 1 (with $T_S = 100$ and $\mu_\lambda = 0.1$). a Evolutionary outcome given antagonism between stimuli and responders ($\phi = f(\pi) = 1 - \pi$). b Evolutionary outcome given mutualism between stimuli and responders ($\phi = f(\pi) = \pi$).

have generated a relatively constant stimulus landscape when there are shared interests. The Goldilocks principle predicts that when the environment of learners is constant, the learning process should be usurped by a fixed response. And indeed a relatively stereotyped response (at birth) evolves under mutualism (Figure 2b). Naïvely, we might thus expect slow or no learning (low λ) given mutualism and more rapid learning (higher λ) given antagonism. It turns out that we find just the opposite (Figure 3b). Even though costly, fast learning is actually a very good way to catch a static target (see Papaj [1994] for a discussion of this point). When dealing with mutualistic stimuli, the responders *produce* this constancy.

We can explore this finding further by running some additional mutualism simulations. For these additional runs, we fix the speed of learning at a low value and do not allow it to mutate ($\mu_\lambda = 0$). We start our responders with a uniform initial behavioral distribution and track responder evolution across generations. We find that the time to evolve the characteristic skewed response of Figure 2b (where one behavior has at least 75% of the weight at birth) is *smaller* when learning is fixed than the case where mutation in learning speed is permitted

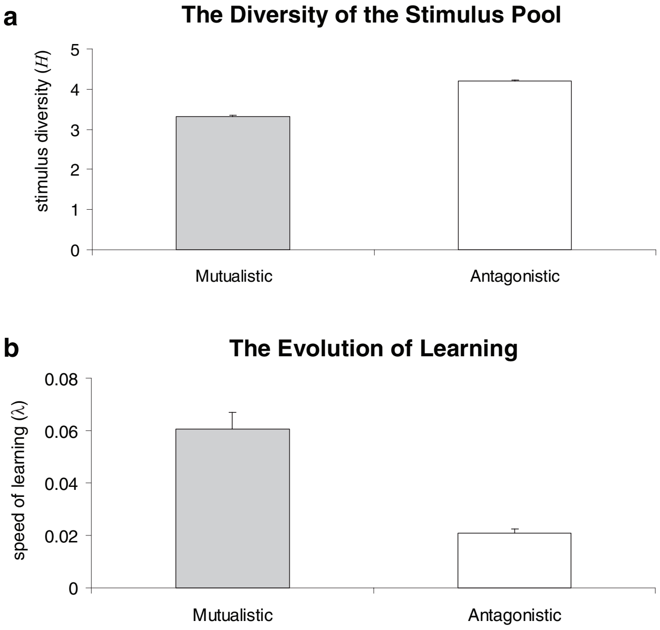

**Figure 3.**
Diversity of stimuli and responder learning evolution. Model parameters are as listed in Table 1 (with $T_S = 100$ and $\mu_\lambda = 0.1$). The "antagonistic" runs assume the payoff relationship $\phi = f(\pi) = 1 - \pi$ while the "mutualistic" runs assume the payoff relationship $\phi = f(\pi) = \pi$. a The Shannon index of diversity ($H$) of stimuli is calculated for each of 10 simulation runs at generation 2500. The average across 10 runs is shown (and error bars give the standard error of the mean). b The average evolved learning parameter (λ) of each of 10 simulation runs is recorded at generation 2500. The average (of mean λ) across 10 runs is shown (and error bars give the standard error of the mean).

(see Figure 4). That is, learning actually *slows* this adaptation process.

The "Baldwin effect" refers to cases in which a trait that is originally ontogenetically acquired, e.g., through learning or phenotypic plasticity, becomes "genetically fixed" (see Baldwin [1896]; Hinton and Nowlan [1987]; Ancel [1999]). Here we see that our responder can achieve high fitness through learning before in-born tendencies in behavior evolve. Responders eventually evolve behavioral tendencies that resemble the skewed distributions produced by learning; thus, there is an element of Baldwin's effect in our model. However, unlike some learning models (e.g., Hinton and Nowlan 1987), we do not observe a "Baldwin Expediting Effect" (see Ancel [2000]) in which learning speeds up the evolutionary attainment of an optimal genotype (in our case a highly skewed behavioral distribution). Learning is favored in a constant world, but it slows "innate" behavioral adaptation in the process. In the case of stimulus–responder mutualism, it is the responders themselves that contribute to the very basis of this constancy.

Given antagonism between the stimulus and responder, the responders construct a diverse stimulus pool (see Figure 3a). As a consequence, any responder will interact with functionally different members of the stimulus population across its lifetime. That is, the payoff for using behavior *i* at

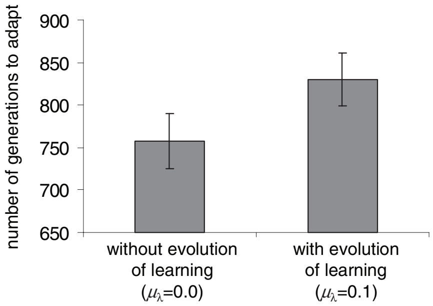

**Figure 4.**
The role of evolvable learning in the time to adapt. Model parameters are as listed in Table 1 (with $T_S = 100$) and all runs assume a mutualism between stimuli and responders ($\phi = f(\pi) = \pi$). Here we measure the time (in generations) for one behavior to evolve 75% of the weight (in the average initial behavioral probability distribution of responders). As shown in Figure 1, mutualistic relations between responders and stimuli favor an initial behavioral distribution where one behavior has nearly all the weight. We run 50 separate simulations under a case where learning cannot evolve ($\mu_\lambda = 0.0$) and where learning can evolve ($\mu_\lambda = 0.1$). In all 100 simulations, the value of λ at the first generation is set at 0.01 (a relatively low value, see Figure 3b) and the initial behavioral distribution at the first generation is set to be uniform. The average time to adapt across 50 runs is shown (and error bars give the standard error of the mean).

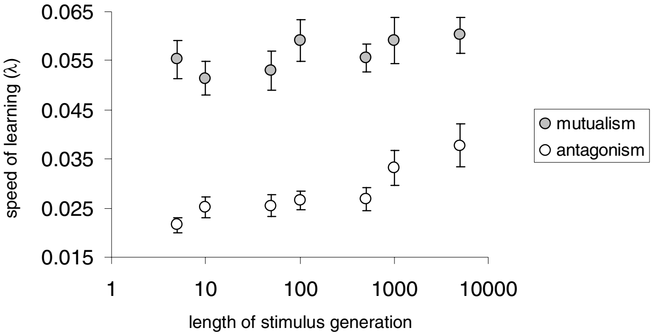

**Figure 5.**
The role of stimulus generation time on the evolution of learning. Model parameters are as listed in Table 1 (with $\mu_\lambda = 0.1$). The "antagonism" runs assume the payoff relationship $\phi = f(\pi) = 1 - \pi$ while the "mutualism" runs assume the payoff relationship $\phi = f(\pi) = \pi$. The average value of the learning parameter (λ) is recorded at generation 2500. Different stimulus generation times are simulated (ranging from 5 time steps to 5000 time steps), while the responder generation time is held constant at 100 time steps. For each stimulus generation time value, 10 simulations are executed. The average (of mean λ) across 10 runs is shown (and error bars give the standard error of the mean).

time step *t* can be different from its payoff at time step *t*+1 if the members of the stimulus population experienced at these time points differ in their evaluation functions. By making its stimulus pool more heterogeneous, the responder makes the learning process less *reliable*. In such a case, responders that reweight their behavioral distribution dramatically after interacting with a rewarding stimulus (responders with a high λ) are selected against.

The above intuition about the cognitive effects of responder-induced variability in stimuli can be checked in another way. While any responder may interact with a heterogeneous set of stimuli across its lifetime, the *pool* of stimuli only changes when the stimuli reproduce. Thus, we can explore a separate contributor to stimulus reliability: the stimulus generation time. Keeping the responder generation time ($T_R$) constant, we vary the stimulus generation time ($T_S$) from cases where many stimulus generations occur with every responder generation ($T_S \ll T_R$) to cases where only a fraction of a stimulus generation occurs with every responder generation ($T_S \gg T_R$). In Figure 5 we see that the evolved speed of learning tends to increase with the $T_S/T_R$ ratio. When the stimulus generations are short, quick learners get abused *within their lifetimes* by a new generation of exploiting mutant stimuli. This tends to select against fast learning when the $T_S/T_R$ ratio is small. When the responder population has several generations with the same stimulus pool (the $T_S/T_R$ ratio is large), faster learning can be selected.

According to the Goldilocks principle, we might expect *intermediate* rates or changeability (e.g., intermediate stimuli diversity or intermediate stimulus generation times) to favor the strongest forms of learning. Instead we find that stronger forms of learning continually evolve as the changeability of the stimulus landscape drops (as stimulus diversity decreases and as stimulus generation time increases). As before, we see that the Goldilocks principle paints an incomplete picture of cognitive evolution in the presence of niche construction.

## Discussion

In an elegant experiment, Mery and Kawecki (2002) explored evolution of learning in fruit flies in real time. Within each generation in their treatment, they paired one of two fruit flavors with an aversive chemical cue (quinine) during a "learning trial" and then removed this cue during subsequent "selection trials" in which flies could choose to oviposit in agar with either fruit flavor. They picked eggs from the flavor not previously associated with quinine for the next generation. The flavor paired with quinine alternated between generations and control populations lacked quinine altogether. These authors found that experimental flies evolved preferential egg laying in the flavor not paired with quinine. Specifically, they found that under the experimental treatment, flies evolved a faster rate of learning, a longer memory, and a pronounced conditioned response in avoiding a flavor paired with an aversive cue.

In this remarkable experiment, the stimulus (quinine) was unaffected by the behavior of the flies. Rather, the experimenters completely controlled the context and reliability of this cue. Thus, cognitive evolution in this experiment does *not* involve the type of niche construction considered above. Indeed, general exploration of cognition within the psychological sciences has been devoid of niche construction. For instance, as rats learn to solve a maze (Tolman 1948) they do not alter the topology of the maze itself, and as a cat discovers a way out of its puzzle box (Thorndike [1911] 2000) it does not alter the state of the box for the next trial. In all of these cases, the world of the learning organism is taken to be (or experimentally forced to be) unaffected by the learner's actions.

Similarly, in the theoretical literature on cognitive evolution, the environment of the cognitive agent is often taken to be unresponsive to the action of the agent. And then the evolution of different cognitive strategies is taken to be a function of different environmental properties. One reason for these externalist perspectives is quite straightforward: it is simpler to deal with a *one-way* causal arrow flowing from environment to agent. In this view, it is the autonomous properties of the environment that explain the cognitive properties of the agent. Niche construction complicates this simple causal picture by making the properties of the environment partially dependent on properties of the agent. A causal loop is introduced, where an organism is both affected by *and affects* the environment (Lewontin 1978, 1982, 1983; Odling-Smee et al. 1996, 2003; Laland and Sterelny 2006).

Given its inherent complexity, is it worthwhile to introduce niche construction in attempts to understand cognitive evolution? In this article, I have tried to make the case that a consideration of niche construction leads to a new perspective on cognition. Ironically, part of the case for niche construction is made by the externalist models for cognitive evolution. These models have shown that the properties of the environment (e.g., the reliability and heterogeneity of stimuli) will influence the optimal cognitive strategy. Thus, it would seem to follow that if the behaving agent can influence these stimuli properties, then its actions could feed back to alter the evolutionary course of its cognitive strategy. Here I have explored how the nature of the feedback affects cognitive evolution. Positive niche construction (where behavior rewarding to the agent selects for the rewarding stimulus) and negative niche construction (where behavior rewarding to the agent selects against the rewarding stimulus) turn out to have qualitatively different consequences for the evolution of cognition.

One element of cognition affected by the nature of the stimulus–responder relationship is the proclivity of the responding organism toward certain behaviors in response to the stimulus. Given that the above model allows the behavioral preferences at birth to evolve, there is the potential for comments on a classic dichotomy: innate versus learned behavior. In reality, the distinction between these categories is not crisp (see Mameli and Bateson [2006] for a full discussion) and upon closer inspection, potential examples within each category contain elements often associated with the other (e.g., the process of learning can sometimes be relatively rigidly constrained; see Gould and Marler [1987]). Regardless of these issues, I will tentatively make a few comments on the innate-learned dichotomy.

In the context of the model, cases where the responder and stimulus have shared interests (a positive $f$ function), a highly skewed behavioral probability distribution at birth evolves. Thus, the organism is "born ready" to behave in a specific way in response to the stimulus and this tendency is resilient to experience with different behaviors. This situation matches some definitions of "innate behavior" fairly closely (e.g., Tinbergen 1989). It is fairly interesting that many (but, by no means, all) of the so-called *innate* behaviors studied by the early ethologists involved cases where the responder and stimulus possess some common interests. For instance, much of this early literature explores behavior involved in parental care of offspring (e.g., egg-rolling behavior in the gray lag goose and feeding strategies in digger wasps), offspring reaction to a feeding parent (e.g., bill pecking in herring gull chicks and gaping in thrush nestlings), and courtship and mating (e.g., the courtship ceremonies of grayling butterflies, claw-waving courtship in the fiddler crab, and female spawning in response to male "quivering" in three-spined sticklebacks) (see Tinbergen [1989]). Indeed, mate recognition and copulation behavior have been called largely innate (Dukas 1998). In these cases, the responder's stimulus is an offspring, a parent, or a mate, and thus some degree of shared interest is certainly plausible. In the context of the model, the positive niche construction is expected to feed back to favor stereotyped behavior. (Note that when the stimulus and responder are conspecifics, the model would have to be altered slightly—e.g., the explicit incorporation of bi-parental inheritance and the specification of genetic linkage between stimulus properties and responder properties).

On the other hand, cases where the responder and stimulus have antagonistic interests (a negative $f$ function), a flat behavioral probability distribution at birth evolves. Some examples of this type of relationship would include predators and their prey or hosts and their parasites. In these cases, the responder "keeps its options open" and within-generation experience can potentially alter this behavior. In the context of the model, the negative niche construction is expected to promote flexibility.

As our model is extremely simple, the above conclusions should be taken as tentative and there are caveats that deserve attention. The stereotyped behavior that is expected to evolve under mutualistic stimulus–responder relationships assumes that all stimuli possess "good intentions." In fact, a highly stereotyped response to a stimulus may constitute an "Achilles heel" of the responder if a stimulus evolves to take advantage of the in-born behavioral tendency. For instance, female fireflies of the *Photuris* genus prey upon male fireflies of the *Photinus* genus by producing flashing patterns that mimic species-specific mating signals (Lloyd 1965; Michaelidis et al. 2006). For the *Photinus* firefly, this is a case where a "good stimulus" (indicating a mate) has "gone bad" (luring it to a predator). The possibility of opportunistic antagonism may affect the degree to which stereotyped behavior evolves. Also, such pressures may contribute to the complexity of the stimulus and response (e.g., several intricate back-and-forth steps in a courtship ritual) with the idea being that more complex stimulus–response relations are harder to exploit.

On the flip side, there may be cases where fundamentally antagonistic relationships produce rigid behavior. One category of particular note includes escape behavior (Dukas 1998; Tinbergen 1989). When the cost of a poor behavioral response is death, it is reasonable to expect the prey to preferentially exhibit behavior carrying a positive probability of escape (even if such behavior is otherwise costly). In the context of the model, certain behaviors (escape behaviors) would always map to low *but consistent* payoffs in response to the predatory stimulus, whereas other behaviors could vary depending on the stimulus individual. Note that our model did not fix any responder behaviors to yield consistent payoffs. With this adjustment to the model, there should be selective pressure to evolve escape tendencies (of course, there would always be pressure for the predator to evolve to exploit these tendencies and make the formerly safe behaviors unsafe again). Generally, however, predator–prey relations potentially exhibit asymmetry, in that the benefits and costs of predatory behavior versus escape behavior can be very different. An unsuccessful predatory behavior means a lost meal for a predator, whereas an unsuccessful escape behavior means the prey item becomes a meal. This means that predators may be better able to afford behavioral flexibility. Thus, predators (more often than their prey) may exhibit the open and adaptable behavioral distributions we find in the antagonistic stimulus–responder scenario above.

How is the nature of the stimulus–responder relationship predicted to affect the responder's learning process? Interestingly, the highest rates of learning (the parameter $\lambda$) occur under mutualistic stimulus–responder relationships. In a world where a given behavioral response produces a consistently high payoff (due to positive niche construction), learning that behavior quickly is selected. Under mutualistic relations, an apparent paradox presents itself: the responder evolves not only a highly skewed initial behavioral distribution (an innate response), but also a high rate of learning. However, this paradox is resolved when one sees that the skew in the in-born behavioral distribution and the learning process accomplish the same end: guiding the organism to a behavior that simultaneously rewards it and the stimulus (see Papaj [1994]). Indeed, Mery and Kawecki (2004) found that fruit flies evolved improved learning over many generations where the same behavioral response was constantly favored. Under mutualistic relations, the responder selects for a constant world—specifically, the most rewarding stimuli are selected. This leads to a reliable world, or what Sterelny (2003) calls *transparent*.

Such transparency is the critical feature of Stephens' model for the evolution of learning (Stephens 1991). Stephens suggests that the reliability needed to favor learning must operate within generations. If stimuli are unreliable at this time scale and if learning is costly (as it is in our model), then learning is selected against (Stephens 1991; Kerr and Feldman 2003). In the case of an antagonistic relation between stimuli and responders, the responder is continually selecting for stimuli that do not reward the more likely behaviors. That is, the negative niche construction of the responder makes its world unreliable, or *translucent*, to use Sterelny's (2003) term. This translucency is apparent in the greater diversity of stimuli (each yielding different payoffs for some behaviors; see Figure 3). Interestingly, even with the heterogeneity promoted by antagonistic relations, responders evolve higher rates of learning when a stimulus generational period lasts longer than the responder's generational period (see Figure 5). Thus, even though the responder is promoting a translucent world, the relative transparency is greater when the same set of stimuli persist over several responder generations. This is a type of "reliable heterogeneity" and it favors a higher learning rate compared to the "unreliable heterogeneity" found under short stimulus generation times.

Our model shows that the simple prediction of the Goldilocks principle (learning is favored for intermediate rates of environmental change) needs qualification. It is true that learning is not favored in a completely unreliable world (and indeed, given antagonism between the responder and stimulus, the responder *constructs* such a world). However, even costly learning turns out to be beneficial in a world that is entirely reliable (for a stimulus–responder mutualism, the responder *makes* its world reliable). This is not to say that the standard alternative to learning (innate behavior) does not evolve in a nonchanging environment. Rather, learning may simultaneously evolve alongside in-born behavioral tendencies. By making "learning" and "innate response" mutually exclusive, the Goldilocks principle does not identify the important role for learning in transparent environments. A responder that exhibits strong behavioral tendencies at birth does evolve under mutualistic scenarios; however, increased rates of learning are concomitantly selected in this circumstance.

The model outlined above is very simple and there are many potential extensions. For instance, we could explore cognitive evolution when there are multiple stimulus classes. Imagine that at any time point $t$ within its lifetime, each responder carries a stimulus preference distribution, which could be represented by the vector

$$\mathbf{x}(t) = \langle x_1(t), x_2(t), x_3(t), \ldots, x_s(t) \rangle,$$

where $x_i(t)$ is the probability of responding to an individual from stimulus class $i$ at time $t$. (Note that with a total of $s$ classes, the responder will possess $s$ behavioral probability distributions). We could use a linear operator framework similar to that outlined above to represent stimulus learning. As before, responders would affect the stimulus population through interaction, potentially altering the frequency and character of different stimulus classes. Phenomena such as search image formation by predators (Bond and Kamil 2002) could be addressed under this extended model. Indeed, it would be interesting if negative niche construction favored the evolution of a flat stimulus preference distribution at birth, where learning about the stimulus classes could lead to different preferred stimulus types within the responder's lifetime.

Another extension of the model involves allowing the stimulus–responder relationship to evolve. This could be accomplished by endowing each stimulus individual with its own $f$ function (which controls the nature of the relationship). In such a case, one could explore the conditions favoring the establishment or breakdown of stimulus–responder mutualisms. Preliminary analyses incorporating this extension suggest that mutualistic relationships can establish in the midst of antagonism.

Of course, extensions and improvements to the way the learning process is modeled could be discussed (e.g., the specific representation of memory). However, the goal here is not to produce the most realistic picture of cognition. Rather, this model is meant to abstract certain elements of the cognitive process in order to investigate how niche construction could influence cognitive evolution. As such, the model may have applications outside of neuronal-based learning. For example, the adaptive immune system possesses some of the properties modeled here and certainly involves a strong form of niche construction.

Despite applications to other systems, the central focus of this model is a deeper understanding of the cognitive process underlying animal behavior. We have seen that the nature of the feedback generated by a behaving agent influences the evolution of its behavioral tendencies and the way its experience alters its future behavior. Behavior may be a particularly potent form of niche construction; however, a full consideration of niche construction may shed light on the evolution of behavior itself.

## References

Ancel LW (1999) A quantitative model of the Simpson–Baldwin effect. Journal of Theoretical Biology 196: 197–209.

Ancel LW (2000) Undermining the Baldwin expediting effect: Does phenotypic plasticity accelerate evolution? Theoretical Population Biology 58: 307–319.

Arnold SJ (1978) Evolution of a special class of modifiable behaviors in relation to environmental pattern. American Naturalist 112: 415–427.

Baldwin JM (1896) A new factor in evolution. American Naturalist 30: 441–451.

Bergman A, Feldman MW (1995) On the evolution of learning: Representation of a stochastic environment. Theoretical Population Biology 48: 251–276.

Bond AB, Kamil AC (2002) Visual predators select for crypticity and polymorphism in virtual prey. Nature 415: 609–613.

Boni MF, Feldman MW (2005) Evolution of antibiotic resistance by human and bacterial niche construction. Evolution 59: 477–491.

Bush RR, Mosteller F (1955) Stochastic Models for Learning. New York: Wiley.

Dukas R (1998) Evolutionary ecology of learning. In: Cognitive Ecology: The Evolutionary Ecology of Information Processing and Decision Making (Dukas R, ed), 129–174. Chicago: University of Chicago Press.

Edwards TC (1989) Similarity in the development of foraging mechanics among sibling ospreys. Condor 91: 30–36.

Feldman MW, Aoki K, Kumm J (1996) Individual versus social learning: Evolutionary analysis in a fluctuating environment. Anthropological Science 104: 209–232.

Godfrey-Smith P (1998) Complexity and the Function of Mind in Nature. Cambridge: Cambridge University Press.

Gould JG, Marler P (1987) Learning by instinct. Scientific American 256: 74–85.

Hinton GE, Nowlan SJ (1987) How learning can guide evolution. Complex Systems 1: 495–502.

Johnston TD (1982) Selective costs and benefits in the evolution of learning. Advances in the Study of Behavior 12: 65–106.

Kerr B (2007) The ecological and evolutionary dynamics of model bacteriocin communities. In: Bacteriocins: Ecology and Evolution (Riley MA, Chavan MA, eds), 111–134. Berlin: Springer.

Kerr B, Feldman MW (2003) Carving the cognitive niche: Optimal learning strategies in homogeneous and heterogeneous environments. Journal of Theoretical Biology 220: 169–188.

Kerr B, Riley MA, Feldman MW, Bohannan BJM (2002) Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. Nature 418: 171–174.

Kerr B, Schwilk DW, Bergman A, Feldman MW (1999) Rekindling an old flame: A haploid model for the evolution and impact of flammability in resprouting plants. Evolutionary Ecology Research 1: 807–833.

Laland KN, Odling-Smee FJ, Feldman MW (1996) The evolutionary consequences of niche construction: A theoretical investigation using two-locus theory. Journal of Evolutionary Biology 9: 293–316.

Laland KN, Odling-Smee FJ, Feldman MW (1999) Evolutionary consequences of niche construction and their implications for ecology. Proceedings of the National Academy of Sciences of the USA 96: 10242–10247.

Laland KN, Sterelny K (2006) Seven reasons (not) to neglect niche construction. Evolution 60: 1751–1762.

Lewontin RC (1978) Adaptation. Scientific American 239(3): 156–169.

Lewontin RC (1982) Organism and environment. In: Learning, Development and Culture (Plotkin HC, ed), 151–170. New York: Wiley.

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

Lloyd JE (1965) Aggressive mimicry in Photuris: Firefly femmes fatales. Science 149: 653–654.

Lorenz K, Tinbergen N (1970) Taxis and instinctive behavior pattern in egg-rolling by the Graylag goose. In: Lorenz K, Studies in Animal and Human

Behaviour, vol. 1, 316–350. Cambridge, MA: Harvard University Press. German orig. 1938.

Mameli M, Bateson P (2006) Innateness and the sciences. Biology and Philosophy 21: 155–188.

Mery F, Kawecki TJ (2002) Experimental evolution of learning ability in fruit flies. Proceedings of the National Academy of Sciences of the USA 99: 14274–14279.

Mery F, Kawecki TJ (2004) The effect of learning on experimental evolution of resource preference in Drosophila melanogaster. Evolution 58: 757–767.

Michaelidis CI, Demary KC, Lewis SM (2006) Male courtship signals and female signal assessment in Photinus greeni fireflies. Behavioral Ecology 17: 329–335.

Odling-Smee FJ, Laland KN, Feldman MW (1996) Niche construction. American Naturalist 147: 641–648.

Odling-Smee FJ, Laland KN, Feldman MW (2003) Niche Construction: The Neglected Process in Evolution. Princeton, NJ: Princeton University Press.

Papaj DR (1994) Optimizing learning and its effect on evolutionary change in behavior. In: Behavioral Mechanisms in Evolutionary Ecology (Real LA, ed), 133–153. Chicago: University of Chicago Press.

Plotkin HC, Odling-Smee FJ (1979) Learning, change and evolution: Inquiry into the teleonomy of learning. Advances in the Study of Behavior 10: 1–41.

Schwilk DW (2003) Flammability is a niche construction trait: Canopy architecture affects fire intensity. American Naturalist 162: 725–733.

Schwilk DW, Kerr B (2002) Genetic niche-hiking: An alternative explanation for the evolution of flammability. Oikos 99: 431–442.

Stephens DW (1987) On economically tracking a variable environment. Theoretical Population Biology 32: 15–25.

Stephens DW (1991) Change, regularity, and value in the evolution of animal learning. Behavioral Ecology 2: 77–89.

Sterelny K (2003) Thought in a Hostile World: The Evolution of Human Cognition. Malden, MA: Blackwell.

Thorndike EL (2000) Animal Intelligence: Experimental Studies. New Brunswick: Transaction. Orig. 1911.

Tinbergen N (1989) The Study of Instinct. Oxford: Oxford University Press.

Tolman EC (1948) Cognitive maps in rats and men. Psychological Review 55: 189–208.

Turner JS (2000) The Extended Organism: The Physiology of Animal-Built Structures. Cambridge, MA: Harvard University Press.