

Evolving Behavioural Choice: An Investigation into Herrnstein’s Matching Law

Anil K Seth

Centre for Computational Neuroscience and Robotics
and School of Cognitive and Computing Sciences
University of Sussex
Brighton, BN1 9QG, UK
anils@cogs.susx.ac.uk

Abstract. In 1961, Herrnstein [4] famously observed that many animals *match* the frequency of their response to different stimuli in proportion to the reinforcement obtained from each stimulus type. Since then, a great deal of research has attempted to elucidate the mechanisms underlying this “matching law”, so far without a clear consensus emerging. Here, we take the view that “choice behaviour” is a product of agent, environment, and observer, and that “mechanisms of choice” are therefore not to be located solely within the chooser. A simple model, employing the novel methodology of evolving choice behaviour in a multi-agent system, is used to demonstrate that matching behaviour can occur (in stable environments) without *any* dedicated choice mechanism.

1 Introduction

All behaviour is choice. R.J. Herrnstein (in [14])

In ALife, behavioural choice has been largely synonymous with *action selection*, which can be loosely defined as the problem of choosing what to do, at any given time, in order to further progress towards multiple, time-varying goals, [8]. Action selection is concerned with choice in the context of different behavioural options that relate to *distinct* goals (for example, feeding and drinking), let us call this a type-1 choice scenario. Here, we consider a different situation; where choice operates between different ways of satisfying the *same* goal - let us call this a type-2 choice scenario - an example would be foraging amongst two kinds of plant that differ in nutritive value. It is in such scenarios in biology that we can observe *matching behaviour* when animals allocate the frequency of their responses to different stimuli in proportion to the reinforcement obtained from each stimulus type.

We may distinguish two questions: (1) how (and why) animals display matching at all, and (2) how animals are able to track changing environmental contingencies. The latter is clearly the more complex, implicating learning, and we shall be concentrating on the former; exploring some minimal conditions on the internal mechanisms required to support matching behaviour in *stable* environments.

In Sections 2 & 3 we briefly revise how the biological sciences of ethology, behavioural ecology, and experimental psychology have approached both type-1 and type-2 choice, and note that ALife, in following an ethological precedent, has (so far) concentrated almost uniquely on the former problem of action selection. The *matching law* of Herrnstein [4] is then introduced as an alternative choice paradigm and it is argued that theories in both domains suffer from what will be a theme throughout this paper; that *behaviour* and *mechanism* should not be confused.

In pursuit of a more fruitful understanding of the mechanisms underlying choice, Section 4 describes a multi-agent model in which simple reactive agents are evolved in a type-2 choice environment. A hypothesis derived from Fretwell [2] is put to the test; that matching behaviour is optimal if there is sufficient *competition* for the resources. This model therefore introduces the evolution of multiple agents for choice behaviour.

The results (in Section 5) support Fretwell's intuition, and indicate that matching behaviour (in stable environments) can arise without there being any dedicated mechanism of choice, and, indeed, without any internalisation of distinct behavioural options at all. Sections 6 and 7 discuss these results in terms of how ALife can contribute towards the natural science attempts to understand choice; in particular, how it can make the relationship between behaviour and mechanism the object of study, rather than the legacy of unwarranted assumptions.

2 Biological Approaches to Behavioural Choice

2.1 Ethology

Ethology, through the observation of animal behaviour in natural contexts, seeks to understand the nature of the mechanisms underlying behaviour *and* their evolutionary origins.¹ Ethological accounts of behavioural choice (in type-1 choice scenarios) generally propose internal mechanisms that arbitrate between internalised (and pre-existing) repertoires of behavioural options. The ALife approach to action selection has also relied on this foundation, and has investigated a broad class of arbitration devices, from hierarchies ([12]) to distributed networks ([7]).

In previous work [10] this approach was criticised for committing the category error of assuming that externally observed behaviours must have internal mechanistic correlates. A behaviour is the joint product of an agent, an environment, and an observer; thus the (agent-side) mechanisms underlying the generation of any behaviour should not be assumed to be identical to the behaviour itself. A simple animat model was used in [10] to illustrate that effective action selection could occur without any internalisation of behaviour, and in which choice could

¹ Tinbergen, a founder of the discipline, actually specified *four* components of any complete ethological account of a behavioural pattern; causation, development, survival value, and evolution (from [9]).

be explained just as well in terms of perception as in terms of action. In this paper, a similar model illustrates that the same arguments apply in understanding the mechanistic basis of behavioural matching.

2.2 Behavioural Ecology

Behavioural ecology (BE) differs from ethology in that it is ostensibly *not* concerned with mechanism, but only with the adaptive rationale for observed patterns of behaviour.² It is within this discipline, and within the complementary discipline of experimental psychology (EP), that the investigation of behavioural matching has acquired greatest momentum.

The study of matching in BE has focussed on *foraging*. Krebs & Kacelnik [6] present a discussion of various foraging strategies, and the environmental conditions under which they may do well. A typical foraging problem is a type-2 choice environment, with two types of source providing reinforcement of a single kind (for example, food), but at different rates or with different probabilities. Two particular foraging strategies are of interest here; the zero-one rule (see [11]) and the matching rule [4]. Given a type-2 environment, these rules make conflicting predictions about how foraging agents should behave with regard to the *less profitable* source (both rules agree that agents should maximize their intake of the most profitable). The former predicts that an agent will *either always* take the less profitable source upon encounter *or never* take it. The latter predicts that the agent will take the less profitable source at a rate that matches the difference in profitability between the two source types.

In Section 3 we shall see how these conflicting predictions can be unified into a single adaptationist framework, but for now let us note that although BE claims to be unconcerned with questions of mechanism, it has been argued that there is often an *implicit* tendency to credit the animal with complex computational and representational abilities (see [3]).³ So either BE tells us nothing at all about mechanism, or it may force us into making strong mechanistic assumptions without appropriate justification.

2.3 Experimental Psychology

It is to experimental psychology (EP) that we must turn in order to find more explicit investigations of the mechanisms underlying behavioural choice. In the case of matching, these investigations are primarily concerned with the more complex issue of how animals learn to track changing contingencies, rather than the simpler question of how (and why) they can match at all. It is perhaps because of this that the status of matching in EP remains unclear, with some believing that matching is the product of underlying learning rules (see, for

² Historically, BE derives from the split in ethology between those exclusively interested in adaptation, and those who concentrated on mechanism (see [1]).

³ For those familiar with BE, this relates to the necessary assumption of a *decision variable* in optimal foraging theory.

example [5]), and others arguing that matching is *itself* the rule by which animals determine their responses, and not a (mere) description of the results (see [14]).

The confusion surrounding the nature of matching may also stem from both the frequent lack of connection between the contrived environments of many EP experiments and ecologically plausible situations⁴, and the common feature of mechanistic EP theories of the concept of *response strength*, which is essentially another manifestation of an internal behavioural correlate.⁵ EP therefore understates the environment both with respect to the behavioural problem (by encouraging arbitrarily constructed environments) and with respect to the mechanistic solution (by relying on internal behavioural correlates). Indeed, given the co-dependence of behaviour on agent and environment, the former inevitably leads to the latter.

Although the emphasis on learning in EP precludes direct comparisons being drawn with the present work, a bridge can be built by considering that the question of learning may productively *follow* an understanding of matching in stable environments, and that the problematic (and complex) mechanistic proposals prevalent in EP (see [14]) can be largely attributed to the confusion between behaviour and mechanism described above.

3 The Matching Law and the Ideal Free Distribution

In Krebs & Kacelnik [6] we find the following definition: “the matching law states that the animal allocates its behaviour in proportion to the rewards it has obtained from them.” More formally:

$$\log\left(\frac{F_A}{F_B}\right) = \log(k) + b\left(\log\left(\frac{r_A}{r_B}\right)\right) \quad (1)$$

where F_A and F_B are the response frequencies to alternatives A and B , and r_A and r_B are the attained rates of reinforcement from the two alternatives, with k and b as scaling parameters. The accepted fact that animals (and possibly humans) often behave according to this rule in behavioural choice situations, both inside and outside the laboratory, is undoubtedly of importance and demands explanation.

In 1972, Fretwell [2] introduced the concept of the *Ideal Free Distribution* (IFD). Given a distribution of resources of different qualities, and a population of (uniformly capable) foragers, the IFD describes the distribution of foragers

⁴ A popular methodology in EP involves creating experimentally contrived abstractions of foraging environments, in which different “choices” are directly presented to animals, with these choices being reinforced according to experimenter-defined “reinforcement schedules”, which can be quite complex (see [14] for a comprehensive review). Note that this methodology is not particular to ‘behaviourism’.

⁵ EP theories often focus on how the “response strengths” are modified, and on how they, in turn, determine the emitted response. But the “response strength” concept is a prime example of a problematic *intervening variable*, crucial to the theoretical coherence of EP, but empirically unsubstantiable from within it.

such that they all do equally well, regardless of the local resource quality, and such that no-one can profit by moving elsewhere. The central intuition is that the high quality resources will tend to become overcrowded, and Fretwell argued that the only foraging strategy that is evolutionarily stable is that foragers match the relative frequency of their choices to the relative qualities of their options, in other words, that they follow the matching law. Of course, in ecological contexts without this element of competition, the optimal behaviour may well be to concentrate exclusively on the higher quality source, in other words to follow the zero-one rule. For my money, this explanation of matching behaviour is satisfying as far as it goes; a multi-forager, limited resource environment is indeed ecologically plausible. But no claims are made about the *mechanisms* involved in this choice behaviour.

It is here that the current paper can contribute empirically, in playing out the above intuition. A model is described which employs artificial evolution to design internal mechanisms, for multiple agents, in a variety of choice scenarios. We are thus able to combine the mechanistic insight sought in EP, with the ecologically motivated adaptationist hypotheses of BE.

4 The Multi-Agent Model

The essence of the model is a type-2 choice environment with multiple foraging agents. A GA is used to evolve the parameters of a reactive feedforward neural network with a fitness function requiring efficient foraging. The hypothesis under test is simple: as the number of foraging agents increases, the behaviour of each agent should approximate that predicted by the matching law. Following the evaluation of this adaptationist hypothesis, the evolved mechanisms are analysed and the relevance of the concept of a “mechanism of matching” is questioned.

4.1 Agent and Environment

The agent(s) exist in a spatially continuous world (but with discrete time steps) containing 4 of each of 2 types of source, *A* and *B*. Each source type can potentially fully replenish the internal battery of the agent (initial level 200), which otherwise diminishes at a steady rate (1 per time step). Source type *A* *always* replenishes the battery ($r_A = 100\%$), but the replenishment probability r_B can be explicitly set at the beginning of each experiment. If the battery reaches 0 the agent dies, and trials terminate after a maximum of 800 time steps.

After an agent has visited a source, the source disappears and reappears randomly in another part of the world. All sources appear in a limited area (200 by 200 units) and the positions of all objects (sources *and* agents) are initialised with a minimum inter-object spacing of 40 units, a condition also adhered to when sources reappear following consumption. Each source is 16 units in radius (5 for each agent), and at full speed an agent covers 2.8 units per time step.

The agents possess 5 sensors, 4 of which are tuned to the two source types (in 2 left/right pairs), and one of which reflects the battery level. The source sensors

respond to the distance from the nearest source of each type to the agent, with each sensor ranging linearly from 100 (at the source) to 0 (≥ 200 units away). If the source is to the left of the agent, the relevant left sensor will respond with 20% greater activation (and *vice-versa* if the object is to the right).

4.2 Genotype and Phenotype

The internal architecture of the agent comprises a simple feedforward reactive network, fully interconnected between layers, but with no internal recurrency within layers. There are 5 inputs feeding through to a 3 unit hidden layer, and then to a 2 unit output layer. The input units linearly scale the sensor values to range from 0 to 1, and all weights range from -1 to 1. Each neuron in the hidden and motor layers applies a sigmoid transfer function to the sum of its inputs (plus a threshold value), with the outputs ranging from 0 to 1. The outputs are scaled to range from -10 to 10 to set the wheel speeds.

These weights and thresholds are specified by 26 real numbers on the 28 number long genotype. The remaining two numbers specify something less orthodox; how well the agent is able to discriminate between the two source types.

Each pair of sensors (S_i), as well as having an associated source type ($i \in \{A, B\}$), also has an associated discriminability value α , ($0 < \alpha < 1$). If both α 's are 0, then both pairs of sensors will behave identically, responding to the nearest source, regardless of type. As either α increases, the associated sensors (S_i) are more likely to respond selectively to their particular source type, i . In detail, this is what happens. First, the distance (from the agent) to the nearest source of each type is calculated (D_i) from which the distance to the nearest source of *any* type, (D_N), is easily derived. Second, for each sensor pair (S_i), if $\mathcal{R} < \alpha$ (where \mathcal{R} is a random number between 0 and 1), then S_i responds to D_i , otherwise to D_N . Once the agent has perceived a source using this procedure, it will continue to perceive it in the same way until a new nearest source comes into range, at which point the procedure repeats.

This additional element of discriminability is admittedly unusual for a simple feedforward network; the rationale for including it derived from a desire to allow the agents the ability to determine (to some extent) their own input, to evaluate the possibility that mechanisms underlying choice can be based on perception as well as action.

4.3 GA and Fitness Function

A distributed GA⁶ was used to evolve population of genotypes in each of 4 conditions in two different models; one model involving a single foraging agent (the *S* model), and another involving 3 agents (the *M* model).⁷ In the *S* model, an incremental fitness function, summed over each time step as $\frac{\mathcal{B}}{200}$ (\mathcal{B} being the

⁶ Population 64, mutation rate 0.04 per bit, crossover rate 0.5. Runs of 1000 generations took about 10 hours on a single user Sun 166MHz SparcStation.

⁷ In the *M* model the agents could not directly perceive each other.

battery level), and averaged over 5 separate trials, rewarded agents that lived long and foraged effectively. In the M model, each genotype gave rise to 3 genetic clones, and the fitness of the genotype was assessed (as above) by following the performance of just *one* of these agents (selected randomly) under the rationale of avoiding direct reinforcement of behaviour beneficial to the group as a whole.

5 Results

Four conditions were first evolved and then tested in both the S and M models, $r_B \in \{100\%, 66\%, 33\%, 0\%\}$, ($r_A = 100\%$ in all conditions). Two were control conditions; $r_B = 100\%$ and $r_B = 0\%$. In the former, source types A and B are functionally identical, and we should expect foragers in both models to attend to each equally. In the latter, type B is always worth nothing, and so we would expect foragers in both S and M models to attend exclusively to type A sources. The experimental conditions were $r_B = 66\%$ and $r_B = 33\%$. Here, we predict that the single S forager will *either* continue to attend to A and B equally, *or* will switch to exclusively attending to A . But in the M model, we predict that each forager will attend to B in proportion to the difference in profitability between A and B .

Fit agents generally evolved after about 500 generations in each condition (in both models), but in each case the population was left until 1000 generations had been completed. The fittest agents from each condition (in both S and M models) were tested 1000 times each, with the average number of visits to A and B sources being recorded. The entire set of evolutions (and analyses) was repeated 10 times to obtain overall averages. Fig 1 illustrates that the above predictions were indeed borne out; in the S model the zero-one rule is followed, but in the M model matching behaviour is observed.

A second analysis, using the same evolved agents, revealed very similar results (fig 2). This time, instead of testing in the same environment as that present during evolution, a parallel with EP was drawn by testing the fittest agent from each condition in the contrived environment of a forced-choice discrimination test. Here, each agent was placed equidistant from a single A source and a single B (and no other sources were present). The trial was stopped as soon as one or other of the sources had been visited, and again each agent was tested 1000 times. Note that these tests always involved a *single* agent, even if evolution had occurred in a multi-agent environment.⁸

Sensitivity to parameter settings was investigated by re-evolving all 8 conditions with either 3 or 5 sources of each type in the environment (10 complete re-evolutions in each case). Analysis was performed as before, and fig 3 illustrates, again, qualitatively similar results to fig 1.

Turning briefly to the evolved mechanisms (of the original agents, not those re-evolved to check parameter sensitivity), the most immediate observation is

⁸ This test environment was meant to bear some similarity to the “skinner-box” experiments widespread in EP.

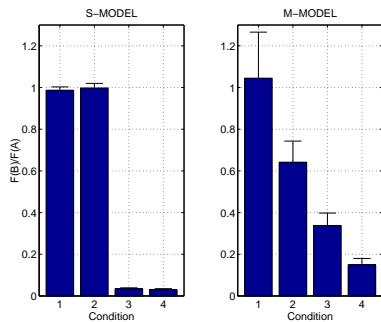


Fig. 1. Same Environment Testing




Fig. 2. Forced Choice Testing

Conditions are 1: $r_B = 100\%$, 2: $r_B = 66\%$, 3: $r_B = 33\%$, 4: $r_B = 0\%$. These graphs show the average (and standard deviation) rate of response to B (F_B) as a fraction of rate of response to A, (F_A) over 10 evolutionary runs in each condition. From each run, a single value was obtained by testing the fittest agent 1000 times in either (fig 1) the same environment as evolution, or (fig 2) in a forced-choice environment. The evolution of matching behaviour is observed in both M-model tests, and the evolution of zero-one behaviour is observed in both S-model tests.

that there is a strong correlation between the difference in response to A and B, and the degree to which the agents can discriminate between the sources (fig 4).

6 Discussion

The above results clearly indicate that matching behaviour can evolve in multi-forager, stable, limited resource environments, and that agents with the same fundamental architecture evolve a zero-one foraging strategy if competition is absent. Four immediate observations can be made. First, the results are robust; neither testing in a forced-choice environment, nor re-evolving with different source densities altered their pattern. Second, it is evident that nothing more than a simple reactive neural network is required on the part of the agent, under the conditions of these models, for either matching or zero-one behaviour (although these agents are unable to track changing contingencies). Third, choice behaviour is apparently mediated as much by perception as by action (fig 4).⁹ Fourth, even in the conditions in which agents appear to be ignoring source type B (when $r_B = 0\%$ in both the S and M models), removal of B sources afflicts their performance (mean fitness difference 2.11%, statistically significant using paired-sample t-test with $t = 2.95$, $df = 18$, $p < 0.01$). Here the B sources are clearly influencing behaviour, even though the agents are never “choosing” B.

⁹ An interesting side issue here is that adaptive behaviour often depends on the sensors presenting less than perfect discriminability. This runs counter to the common intuition that the more accurately sensors reflect the external world, the better.




Fig. 3. The distinction between matching and zero-one behaviour remains clear with altered source densities. The two leftmost graphs report results with 3 sources, and the two rightmost, 5 sources. The graphs are to be interpreted as in fig 1. All testing was in the same environment as evolution, and 10 evolutionary runs were performed in each condition.




Fig. 4. For each fittest agent in each condition (a total of 80 data points) the ability of the agent to discriminate between A and B sources is plotted against the average response pattern (over 1000 trials) to the sources. There is a clear correlation.

These observations suggest that the idea of an explicit “choice mechanism” (or worse, a specific “mechanism of matching”) supervening on an internalised repertoire of behavioural options is inappropriate for these evolved agents. However, this is *not* to say that the differences between the mechanisms are irrelevant to the differences between the behaviours; the forced-choice test results (fig 2) indicate that agents evolved in *M* (multi-agent) models still behave according to the matching law even when tested *without* conspecifics.

If stable matching behaviour can arise without there being a dedicated “mechanism of matching”, it is natural to ask whether the importance of the matching “law” itself has been overstressed. One possibility is that the emphasis placed on matching derives from the fact that most (if not all) studies of matching (in biology) have taken place in variants of type-2 choice environments. However as Williams [14] argues; the way in which *other choice alternatives* exert their effects “is a vital theoretical issue that remains to be clarified.” Current work is addressing this issue by evolving choice behaviour in environments with both type-1 and type-2 characteristics (thereby also permitting the joint examination of the problems of matching and action selection.) Future work may also address the evolution of learning mechanisms that allow agents to track changing contingencies. It is perhaps here that parallels with EP might be more in evidence, and that the notion of a “mechanism of matching” may find greater application.

It appears that the methodology espoused here provides a promising way of addressing the difficulties faced by both EP and BE in how they deal with

mechanism (see Section 2). By making the distinction between behaviour and mechanism fully explicit, and the object of study rather than assumption, we can avoid the pitfalls of either (1) searching for mechanistic correlates that underlie behaviours in arbitrary and contrived environments in EP, or (2) relying on the hidden mechanistic assumptions necessary for the construction of predictive models in BE. But there is an obvious pitfall for the ALife approach as well; the currency of minimal models and simple environments may bear very little relation to the complexities of their biological counterparts. The temptation to use ALife to attempt to displace biological knowledge must therefore be resisted. It is better that the insights from such simple modelling are used to engender shifts in conceptual stance of benefit to all intellectual parties.¹⁰

A word of caution is therefore necessary. I do not suggest that animals (or humans) choose in the way that these evolved agents choose. Nor do I ignore the fact that both humans and animals display matching behaviour in a much wider range of situations than has been investigated here, and are clearly able to rapidly modify their behaviour in response to continuously changing reinforcement contingencies. All I suggest is that the assumption that observed matching behaviour need be supported by a dedicated “mechanism of matching” supervening on internalised behavioural correlates (or response strengths) should not be entirely trusted.

7 Conclusions

Marian Stamp Dawkins [1] has maintained for some time that a satisfying framework for understanding animal behaviour will require a new appreciation of Tinbergen’s ideal that both adaptation and mechanism be considered jointly. As she says, the behavioural ecologists must start talking to the neurobiologists (or, indeed, the experimental psychologists). I would like to conclude by suggesting that ALife can encourage this dialogue; in the present paper we have seen an exploration of a BE phenomenon, in which an adaptationist hypothesis (Fretwell’s IFD) has been cashed out in mechanistic terms. We have seen that the search for a “mechanism of matching” is easily wrongfooted, given that choice (or, indeed, *any* behaviour) resides not solely inside the chooser, but in the joint activity of the chooser, its environment, and the observer. We have seen that matching behaviour in stable environments need not require anything more than a simple reactive mechanism.

This work also benefits ALife itself, in three ways; by highlighting the potential importance of perception in choice (a theme further developed in [10]), by demonstrating the evolution of choice behaviour in a multi-agent system, and not least by drawing attention to a rich body of behavioural choice phenomena to which the current paper serves as the briefest of introductions.

¹⁰ There is of course a different tradition in ALife that *is* concerned with developing models of stringent biological plausibility, from which biological hypotheses can be directly assessed (see [13] for example.)

To return to Herrnstein's aphorism that "all behaviour is choice". Well, one could equally say that *no* behaviour is choice; behaviour is just observed ongoing agent-environment activity. We, as observers, demarcate and label portions of this stream of activity, and sometimes call the boundaries choices. The foraging agents in this paper do not employ any mechanisms of choice; of course how animals may work remains an open question.

Thanks to Matt Quinn, Andy Philippides, Emmet Spier, Tom Smith, to my supervisors Phil Husbands and Hilary Buxton, to J. Peck for introducing me to the work of Herrnstein, and to the EPSRC for financial support (award no. 96308700).

References

1. M.S. Dawkins. *Unravelling animal behaviour*. Longman, London, 1995. 2nd edition.
2. S.D. Fretwell. *Populations in seasonal environments*. Princeton University Press, Princeton, 1972.
3. C.R. Gallistel, A.L. Brown, S. Carey, R. Gelman, and F.C. Keil. Lessons from animal learning for the study of cognitive development. In S. Carey and R. Gelman, editors, *The epigenesis of mind*, pages 3–37. Lawrence Erlbaum, Hillsdale, NJ, 1991.
4. R.J. Herrnstein. Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4:267–272, 1961.
5. J.M. Hinson and J.E.R. Staddon. Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior*, 39:25–47, 1983.
6. J.R. Krebs and A. Kacelnik. Decision making. In J.R. Krebs and N.B. Davies, editors, *Behavioural ecology*, pages 105–137. Blackwell Scientific Publishers, Oxford, 1991. 3rd edition.
7. P. Maes. A bottom-up mechanism for behaviour-selection in an artificial creature. In J.-A. Meyer and S.W. Wilson, editors, *From animals to animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior*, pages 238–246, Cambridge, MA, 1990. MIT Press.
8. P. Maes. Modelling adaptive autonomous agents. *Artificial Life*, 1:135–162, 1994.
9. D.J. McFarland. *Anim. Behav. (3rd ed.)*. Longman, 1999.
10. A.K. Seth. Evolving action selection and selective attention without actions, attention, or selection. In R. Pfeifer, B. Blumberg, J.A. Meyer, and S.W. Wilson, editors, *From animals to animats 5: Proceedings of the Fifth International Conference on the Simulation of Adaptive Behaviour*, pages 139–147, Cambridge, MA, sep 1998. MIT Press.
11. D.W. Stephens and J.R. Krebs. *Foraging theory*. Monographs in behaviour and ecology. Princeton University Press, Princeton, 1986.
12. T. Tyrrell. An evaluation of Maes' bottom-up mechanism for behavior selection. *Adaptive Behavior*, 2(4):307–348, 1994.
13. B. Webb. Robotic experiments in cricket phonotaxis. In D. Cliff, P. Husbands, J. Meyer, and S. Wilson, editors, *From animals to animats 3: Proceedings of the 3rd international conference on the simulation of adaptive behaviour.*, pages ???–???, Cambridge, Mass., 1994. MIT Press.
14. B.A. Williams. Reinforcement and choice. In N.J. Mackintosh, editor, *Animal learning and cognition*, pages 81–108. Academic Press, New York, 1994.