

Competitive Foraging, Decision Making, and the Ecological Rationality of the Matching Law

Anil K Seth^{1,2}

¹The Neurosciences Institute, 10640 John Jay Hopkins Drive, San Diego, CA 92121, USA

²CCNR/COGS, University of Sussex, Brighton, BN1 9QH, UK

seth@nsi.edu

Abstract

The matching law describes how *individual* foragers often allocate their choices, occasionally suboptimally, in experimental situations. The ‘ideal free distribution’ predicts how *groups* of foraging agents should distribute themselves, optimally, over patchy environments. This paper explores the possibility that a single behavioural heuristic can account for both phenomena, allowing the potential suboptimality of matching to be understood in terms of adaptation to a group context. Two simple heuristics are compared, ϵ -sampling and ω -sampling: the latter is successful in both cases, but contrary to prior claims in the literature the former is successful in neither. These results emphasise the importance of multiple environmental value estimates in effective decision making.

1 Introduction

The principle of ‘ecological rationality’ holds that cognitive mechanisms are best understood as fitting the demands and structure of particular environmental niches, as opposed to the classical view of cognition approximating a ‘Laplacean superintelligence’ for achieving “general purpose, optimal performance in any situation” (Bullock & Todd, 1999, p.3). An implication of ecological rationality is that cognitive mechanisms operating outside of their proper niche may deliver *irrational* behaviour, in much the same way that a fish out of water is disadvantaged with respect to breathing.

This paper focuses on Herrnstein’s (1961) ‘matching law’, which describes how *individual* foragers allocate their choices in many experimental situations. Matching behaviour is very widespread, characterising much decision making activity in both humans and animals, and importantly, whilst often optimal, it is not always so. Here, we explore the idea that matching can be accounted for by an ecologically rational decision making heuristic well adapted for decision making in a *group* context, from this perspective the occasional suboptimality

(or irrationality) of matching would be a consequence of shifting from a group to an isolated individual environment. A broader motivation is to provide an alternative, empirically grounded perspective on the problem of ‘action selection’ (see Tyrrell, 1993, for example), a term which is often used synonymously with decision making in the field of adaptive behaviour.

The group context we consider is the ‘ideal free distribution’ (IFD), which describes the equilibrium (and optimal) distribution of foragers over an environment with patchy resource distribution (Fretwell, 1972). Even though the IFD and the matching law derive from different disciplines (the former from behavioural ecology, the latter from psychology), they present many similarities: (1) the matching law is to do with individual choice and the IFD is to do with its collective consequences; (2) they display striking mathematical congruency (Gray, 1994; Baum & Kraft, 1998, see also sections 2 and 3); and (3) perhaps most importantly, the laboratory environments employed by psychologists are often interpreted as abstractions of natural foraging environments (Dallery & Baum, 1991; Shettleworth, 1988). The suggestion that a single behavioural dynamic may underlie both phenomena is therefore not new (Houston, 1986; Gallistel et al., 1991; Thuisjman et al., 1995; Seth, 1999). The present contribution is primarily to explicitly assess candidate heuristics, to demonstrate success in one case, and to repudiate equivalent claims of success in another (Thuisjman et al., 1995).

Two simple heuristics are compared, ϵ -sampling and ω -sampling (fully defined later), differing primarily in that the former maintains a single estimate of environment ‘value’ and the latter maintains multiple estimates. Contrary to the claims of Thuisjman et al. it is shown that ϵ -sampling fails to account for either matching or the IFD, but that the novel ω -sampling heuristic is successful in both cases.

2 The ideal free distribution

Given a patchy distribution of resources, the IFD describes the equilibrium distribution of foragers such that no forager can profit by moving elsewhere, regardless of

the local resource quality (Fretwell, 1972); in this state all foragers will obtain equal resources.¹ In order to make specific predictions, IFD models require a way of relating the per forager intake rate W_i (s^{-1}) to both the forager density N_i and the resource availability F_i on each patch i . The following equation is adapted from Sutherland (1983) and Milinski and Parker (1991):

$$W_i = \frac{QF_iF^*}{N_i^m}, \quad (1)$$

in which Q (ms^{-1}) is a measure of patch-independent forager efficiency, $F_i \in [0.0, 1.0]$ (dimensionless) represents the resource fraction in patch i , F^* represents the total resources available, and m (dimensionless) is the interference constant, which is usually taken to vary between 0.0 (no interference) and 1.0 (high interference), where interference is defined as the (more-or-less immediately reversible) decline in intake due to the presence of conspecifics (Goss-Custard, 1980; Sutherland, 1983).² Across two patches A and B , assuming $W_A = W_B$ (the IFD condition):

$$\log \frac{N_A}{N_B} = \frac{1}{m} \log \frac{F_A}{F_B}, \quad (2)$$

this being the ‘generalised habitat matching rule’ of Fagen (1987). Also, taking the total forager number to be $N_T (= N_A + N_B)$, it is possible to predict both N_A and N_B directly (Tregenza, Parker, & Thompson, 1996):

$$N_A = \frac{N_T}{(10^{-c} + 1)}, \quad c = \frac{\log \frac{F_A}{F_B}}{m}. \quad (3)$$

In what follows, we will use equation 3 to assess the ability of ϵ -sampling and ω -sampling to lead populations to the IFD under different levels of interference.

2.1 Resource allocation

Both ϵ -sampling and ω -sampling operate over discrete time intervals, and as such it is possible to interpret the resource level F_i in at least two ways. The first is simply as specifying a resource level that contributes to the forager intake at every time-step. This process of ‘continuous allocation’ (C-allocation) is the usual interpretation in the literature (see, for example, Bernstein, Kacelnik,

¹This last clause requires the assumption that foragers are able to move, without cost, to the patch in which their rewards are maximised.

²Values of m in excess of 1.0 are possible, and can be expected in cases in which prey items can be *lost* (for example, by fleeing) as a result of interference. Note also that the model described here is known in the literature as a ‘standing stock’ or ‘interference’ model for the reason that it assumes relatively constant resources in each patch. Another popular choice is the ‘immediate consumption’ model, which assumes a steady input of immediately consumed resources. With $m = 1.0$ in equation 1, but not otherwise, the two models are equivalent; see Van der Meer and Ens (1997) for further discussion.

& Krebs, 1991), and in this case equation 1 can be used at every time-step, exactly as it is written. The second approach is to understand F_i as as specifying a probability that patch i will yield the fixed resource quantity F^* to each forager at each time-step. Under this process of ‘probabilistic allocation’ (P-allocation), W_i becomes a random variable:

$$W_i = \begin{cases} \frac{QF^*}{N_i^m}, & p(F_i) \\ 0, & p(1 - F_i) \end{cases} \quad (4)$$

The IFD condition of equal intake rates across all patches in this case must apply to *expected* intake rates over many time-steps. We can write:

$$E(W_i) = \left(\frac{QF^*}{N_i^m} \right) F_i.$$

from which the condition $E(W_A) = E(W_B)$ leads to the same generalised habitat matching law described above (equation 2).

No claims are made for the biological relevance of the distinction between C-allocation and P-allocation, it is motivated by analogous resource allocation methods often employed in ‘matching law’ experiments, described below. One possible intuition, however, is that it may reflect a difference between relatively accessible and widespread types of resource (grass, for example), and relatively inaccessible yet potent types of resource (truffles, for example).

3 The matching law

Moving on to the matching law, Krebs and Kacelnik (1991) offer the following definition: “the matching law states that the animal allocates its behaviour to two alternatives in proportion to the rewards it has obtained from them” (p.131). If the proportionality is direct, this is known as ‘strict’ matching (Davison & McCarthy, 1988):

$$\frac{B_A}{B_B} = \frac{R_A}{R_B}, \quad (5)$$

where B_A and B_B represent the rate of response to options A and B , and R_A and R_B represent the resources obtained in each case. The ‘generalised’ matching law (Baum, 1974) includes parameters for bias (b) and sensitivity (s) to account for the departures from strict matching often observed in empirical data:

$$\log \frac{B_A}{B_B} = s \log \frac{R_A}{R_B} + \log(b). \quad (6)$$

Of the similarities between the IFD and matching noted earlier, their mathematical congruence should now be particularly evident (compare equations 2 and 6), but there are also profound differences: whereas habitat matching predictions are normative, the individual

matching law is an observed relation, and whereas habitat matching is expressed in terms of *available* resources (F_i), the individual matching law is expressed in terms of *obtained* resources (R_i). (Notice that F_i can still be used in the context of individual matching even if it is not represented in the matching equations themselves, and indeed it is necessary to do so in order to describe the various ‘schedules of reinforcement’ by which resources are allocated in matching experiments.³)

Psychologists have investigated matching under many different reinforcement schedules, four of which are considered here, each with two options A, B with associated resource availabilities F_A, F_B :

- *Basic*: Each response is rewarded with an amount determined by the relative values of F_A and F_B . Responses are rewarded at every time-step. This is analogous to C-allocation.
- *Concurrent (conc) VR VR*: A variable ratio (VR) schedule indicates that an option must receive a certain number of responses before a reward is given. This number can vary around a mean value, and can therefore be implemented by associating a probability of reward with each option. F_A and F_B are here interpreted as the mean values, so that conc VR VR is analogous to P-allocation.
- *Concurrent (conc) VI VI*: A variable interval (VI) schedule requires that a certain delay elapse after a reward on a given option until that option can be rewarded again. This delay time can vary around a mean, and these means can differ between response options (F_A and F_B are interpreted as the delays).
- *Concurrent (conc) VI VR*: This is a ‘mixed’ schedule in which one choice option is rewarded under a VI schedule, and the other under a VR schedule.

Under both basic and conc VR VR schedules, the general consensus in the literature is that exclusive choice for the most profitable option is observed. There is nothing counterintuitive about this; if repeatedly offered a choice between 80p and 40p, any sensible subject would presumably choose the former 100% of the time, and the same would apply to repeated choices between odds of 3:1 and odds of 5:1. Observations of exclusive choice, although consistent with the matching law, are only trivial instances of its applicability, as such these schedules present relatively undemanding assessments of matching behaviour.

The conc VI VI schedule is more interesting. Unlike basic and conc VR VR, the reward rate can be largely independent of the response rate, such that matching

³The terms ‘reinforcement’ and ‘reward’ are used interchangeably; ‘reinforcement’ is employed only when it helps to maintain consistency with the psychological literature.

to obtained resources can be achieved with a variety of response distributions, including - but not limited to - exclusive choice. Furthermore, under conc VI VI, exclusive choice is no longer optimal (Herrnstein, 1970). Matching to obtained resources under conc VI VI has been observed for both non-humans animals (Davison & McCarthy, 1988) and human subjects (Conger & Killeen, 1974), in all cases *without* exclusive choice.

The final schedule, conc VI VR, also leads to observations of matching to obtained resources, in some cases in the trivial form of exclusive choice, and in other cases non-trivially, depending on the relative productivities of the two component schedules (Herrnstein & Heyman, 1979; Herrnstein & Vaughan, Jr., 1980). The most important feature of this schedule is its relation to the reward maximisation. Unlike all previous schedules, matching to obtained resources (whether trivial or not) is *not* optimal. Conc VI VR therefore enables exploration of the potential suboptimality associated with matching. The consensus in the literature is that matching to obtained resources - not maximisation - is observed under conc VI VR (Herrnstein & Heyman, 1979; Herrnstein, 1997).

4 A Description of the model

With this background in place, we can now turn to the heuristics themselves.

4.1 ϵ -sampling

The idea behind ϵ -sampling is simply that agents stay on a ‘current’ patch, and occasionally ‘sample’ other patches, switching if and only if the ‘sampled’ patch is better.

More formally, given two alternatives A and B , an ϵ -sampler initially selects A or B at random. At each subsequent time interval, it abides by its choice with probability $(1 - \epsilon)$, and samples with probability ϵ , remaining with the sampled option (with probability $1 - \epsilon$) if the reward from this option exceeds a ‘critical level’ (E), which is a dynamic estimate of the ‘value’ of the environment in which more recent rewards are more strongly represented to a degree specified by an *adaptation rate* γ . The operational definition of ϵ -sampling given below is from Thuisjman et al. (1995):

Definition 1 Let $\gamma, \epsilon \in (0, 1)$, let $M(t) \in A, B$ represent the option selected and let $r(t)$ be the resources obtained at time $t \in \{1, 2, 3 \dots\}$. Define $E(1) = 0$ and

$$E(t + 1) = \gamma E(t) + (1 - \gamma)r(t)$$

for $(t \geq 1)$. Then $E(t)$ is called the critical level at time t . Let A_ϵ denote the behaviour of choosing A with probability $(1 - \epsilon)$ and B otherwise. Let B_ϵ be defined similarly. The ϵ -sampling strategy is then defined by playing:

at $(t = 1)$ use $A_{0.5}$,
at $(t = 2)$ use $M(1)_\epsilon$,
at $(t > 2)$ use $M(t-1)_\epsilon$ in case $M(t-1) \neq M(t-2)$
and $r(t-1) > E(t-1)$, otherwise use $M(t-2)_\epsilon$.

4.2 ω -sampling

The novel ω -sampling heuristic extends ϵ -sampling by allowing patch switching to be driven directly by estimates of patch value as well as by sampling excursions; ω -samplers also maintain concurrent estimates of *each* (visited) patch, rather than (as for ϵ -sampling) a single estimate of environmental quality as a whole. (The implications of relaxing this strong assumption in relatively complex environments are discussed in section 5.2.)

For a two patch environment, a ω -sampler initially selects A or B at random. At each subsequent time interval, the other option is sampled with probability ϵ , otherwise (with probability $1 - \epsilon$) the estimate of the current selection is compared with that of the unselected option, and switching occurs if the former is the lower of the two. Operationally:

Definition 2 Let $\gamma, \epsilon, M(t), r(t)$ be as in Definition 1, let $E_A(t)$ and $E_B(t)$ represent the estimated values of options A, B , and let $N(t)$ represent the unselected option at time $t \in \{1, 2, 3 \dots\}$. Define $E_A(1) = E_B(1) = 0$. For $(t \geq 1)$ then if $M(t) = A$:

$$E_A(t+1) = \gamma E_A(t) + (1-\gamma)r(t), \quad E_B(t+1) = E_B(t),$$

otherwise (if $M(t) = B$):

$$E_A(t+1) = E_A(t), \quad E_B(t+1) = \gamma E_B(t) + (1-\gamma)r(t).$$

Let $\mathcal{R} \in (0, 1)$ be a random number. Let A_ϵ and B_ϵ be as in Definition 1. The ω -sampling strategy is then defined by playing:

at $(t = 1)$ use $A_{0.5}$,
at $(t = 2)$ use $M(1)_\epsilon$,
at $(t > 2)$ if $(\mathcal{R} < \epsilon)$ use $N(t-1)$, else if
 $(E_{M(t-1)} < E_{N(t-1)})$ use $N(t-1)$, otherwise use
 $M(t-1)$.

It would not do to overstate the novelty of ω -sampling. Many similar strategies are described in the theoretical biology literature, and it is certainly comparatively trivial in relation to the many reinforcement learning algorithms described in the computer science literature. What is novel here is application to matching and the IFD, and simplicity in this context can be considered a bonus.⁴

⁴There are particularly evident similarities between ω -sampling and decision rules based on the ‘marginal value theorem’ (Charnov, 1976) which have long been associated with the IFD (but not with matching; see, for example, Bernstein et al., 1991), and which specify switching whenever the gain rate in a given patch is lower than the expected gain rate for the environment as a whole. ω -sampling may be considered a marginal-value rule augmented by (1) patch-specific value estimates and (2) sampling-driven switching.

4.3 Model structure

Both heuristics were explored using individual-based models to assess their performance in three conditions: (1) ability to lead groups of agents to the IFD, (2) matching performance of agents when isolated, and (3) matching performance when embedded in a group (this latter condition included because the few empirical biology papers that consider both matching and the IFD generally consider only embedded individuals).

The first condition involved recording the equilibrium distribution (after 1000 time-steps) of populations of 100 ϵ - and ω -samplers, for each of 9 different resource distributions across two patches A and B . Four separate populations were analysed for each strategy, one for each combination of interference level (1.0 or 0.3) and allocation method (C-allocation or P-allocation). In each case, agents were initially randomly allocated to either A or B . Then, each time-step, the resource obtained by each agent was calculated (equation 1 for C-allocation and equation 4 for P-allocation), the appropriate heuristic applied, and the new agent distributions determined. The final equilibrium distributions were compared with the predictions of the IFD (equation 3).

Isolated individual behaviour was analysed under various reinforcement schedules. For the basic and conc VR VR schedules, single ϵ - and ω -samplers foraged in isolation, under C-allocation or P-allocation respectively, for 1000 time-steps under each of 9 different resource distributions. Conc VI VI was implemented by using F_i to set delay intervals (D_i) such that $D_i = 20(1.0 - F_i) + r$, with $r \in [-2, 2]$ an integer random number. The first response to option i on each evaluation procured the full reward F^* and initialised D_i . Subsequent responses to i went unrewarded until D_i time-steps had elapsed, after which a response would again procure F^* and re-initialise D_i , with the incorporation of r ensuring that the schedule was indeed ‘variable interval’. Conc VI VR was implemented by applying VI to one option (A), and P-allocation to the other (B). As before, isolated agents were allowed to forage for 1000 time-steps under each of 9 different resource distributions.

The final condition involved recording the behaviour of individuals embedded within their respective groups from each of the 4 original populations over the full 1000 time-steps, under each of the 9 resource distributions, comparing their behaviour with the predictions of the individual matching law.

All three conditions were repeated 30 times each, enabling means and standard deviations to be calculated.

4.4 Parameter values

Both heuristics require values to be chosen for ϵ and γ . Rather than relying on arbitrary choice and holding this choice constant across all conditions, as is usually the

case, in this study a genetic algorithm (GA) was used to evolve near-optimal values for each condition (see Appendix A for details). The reasoning behind this is as follows. The objective of comparing heuristic performance over a range of conditions requires some equivalence criteria to be drawn in terms of the parameters ϵ and γ . If a fixed parameter set is chosen, it could be argued that because the parameters themselves are identical in all conditions, any performance differences must be due to inherent strategy properties, however an alternative interpretation is that the arbitrary set may be more appropriate for some conditions than others, and so performance differences may, to some extent, reflect imbalances in parameter suitability rather than inherent strategy properties. An alternative equivalence criteria is that of optimality. Optimal (or near-optimal) parameter values may well vary across conditions, but it can now be asserted that, in each condition, each strategy is performing as well as it possibly can, therefore any performance differences really must reflect inherent strategy properties, and cannot be explained away in terms of parameter (un)suitability. This is the intuition followed in the present study.

Importantly, whilst this methodological point is worth making and has been largely overlooked in the literature, the following results do not depend on it. Identical results (not shown here, see Seth, 2000) were obtained from a control study performed using a fixed parameter set derived from near-optimal values averaged over all conditions ($\bar{\epsilon} = 0.052$ and $\bar{\gamma} = 0.427$). These control results also proved robust to small variations in this mean near-optimal set (Seth, 2000).

5 Results

5.1 ϵ -sampling

Figure 1 compares observed distributions of ϵ -sampling agents to the predictions of the IFD (equation 3). Although in most cases there is a good match, ϵ -sampling agents are unable to find the IFD under P-allocation with $m = 0.3$.

With regard to matching (figure 3), ϵ -sampling agents exhibit exclusive choice (trivial matching) under basic reinforcement, in line with the psychological data (3a,e). Under conc VR VR, however, although they continue to match to available resources, they no longer match to obtained resources, and certainly do not exhibit exclusive choice (3b,f). Performance is no better under conc VI VI or conc VI VR; in both cases there are clear departures from strict matching to obtained resources (3g,h).

Embedded ϵ -sampling agents in most cases match closely to obtained resources (figure 4), although there is some divergence from strict matching when $m = 0.3$ under P-allocation (4h).

Contrary to the claims of Thuisjman et al. (1995),

these results demonstrate that ϵ -sampling can neither reliably lead populations of agents the IFD, nor reliably lead individual agents to match to obtained resources. It is worth asking why Thuisjman et al. reached such different conclusions, and one likely reason is that they considered only a small set of analytically tractable special cases. With respect to the IFD they explored only C-allocation with $m = 1.0$, a condition in which ϵ -sampling does indeed lead populations to the IFD, but there is at least one other condition (representative of many others) in which it does not. They also analysed isolated ϵ -samplers only under the equivalent of conc VR VR, which not a useful way to explore matching since only trivial adherence (exclusive choice) is to be expected. Moreover, ϵ -sampling does not even deliver this, instead leading agents to match to available resources. Unfortunately, this result, which Thuisjman et al. (1995) *did* obtain, was wrongly asserted by them to be consistent with the individual matching law; it appears they simply misunderstood the matching law as pertaining to available resources. Here we have seen that ϵ -sampling matches to obtained resources only under basic reinforcement. It should be stressed that the present results are in agreement with those of Thuisjman et al. (1995) *in those special cases considered by them*. The problem is that these cases are not sufficient for substantiating their claims.

5.2 ω -sampling

Turning to ω -sampling, it is immediately clear that it outperforms ϵ -sampling at least with respect to the IFD. In all 4 conditions ($m = 1.0$ or 0.3 , C-allocation or P-allocation) ω -sampling populations closely fit equation 3 (fig 2).

Matching performance is also improved by ω -sampling. ω -samplers agents exhibit exclusive choice (trivial matching to obtained resources) under basic and conc VR VR (figure 5).⁵ Furthermore, matching to obtained resources is also observed under both conc VI VI and conc VI VR (figure 5g,h). (The slight deviations from strict matching entailed by ω -sampling under conc VI VR are in the *opposite* direction to that expected if agents were maximising reward; see Herrnstein & Heyman, 1979.)

Embedded ω -samplers also reliably match to obtained resources under all 4 test conditions (figure 6). Notice, however, that these observations are not reflected in the relatively accessible (in the field) statistic of matching to *available resources*. The significance of this is that if embedded agents are *not* observed to match to available

⁵Careful inspection of fig 5(a,b) reveals that the exclusive choice of ω -sampling under conc VR VR is not quite as exclusive as it is under basic. Although this deviation is slight, it is interesting to note that similar deviations have also been observed empirically (see, for example, Sutherland & Mackintosh, 1971).




Figure 1: Observed (solid) and predicted (dashed) ϵ -sampling population distributions under 9 different resource distributions (IFD predictions obtained using equation 3). Each observation derives from the mean of 30 distributions, standard deviations are shown. Each abscissa represents F_A and each ordinate represents the percentage of agents on patch A. Four conditions are shown, defined by all combinations of interference level (1.0 or 0.3), and C-allocation or P-allocation (CA or PA).

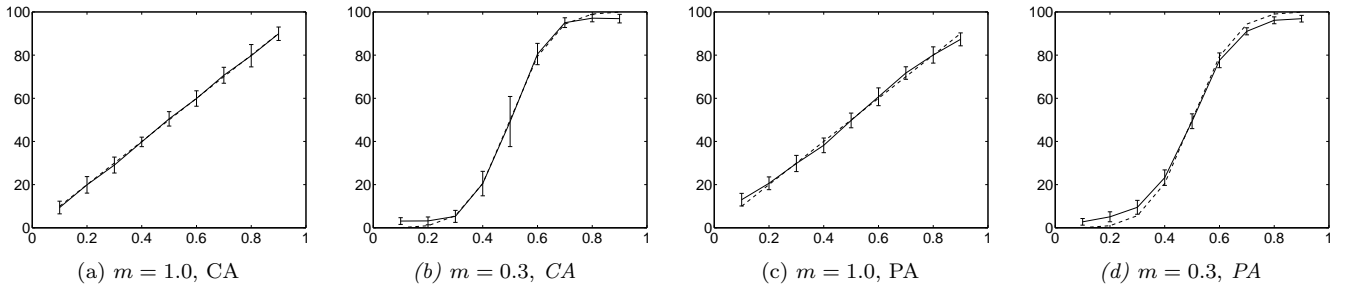


Figure 2: Observed (solid) and predicted (dashed) ω -sampling population distributions under ω -sampling, to be interpreted as figure 1.

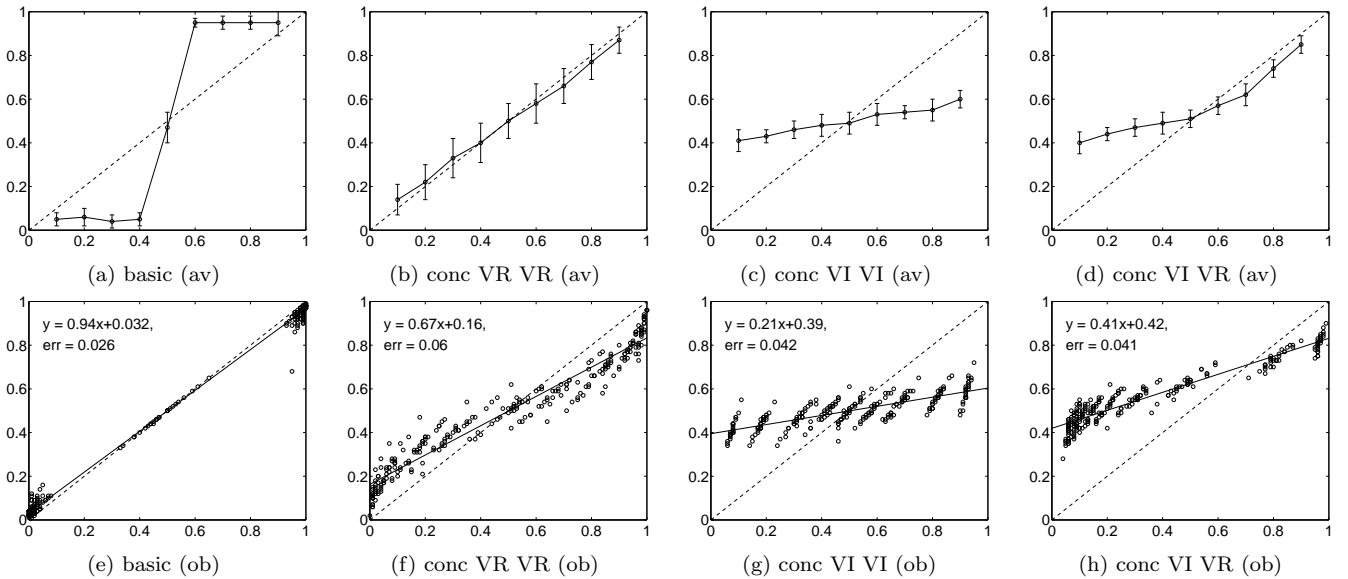


Figure 3: Matching behaviour of isolated ϵ -sampling agents. Data for each plot is collected from 30 analyses at each of 9 values of F_A , with dashed lines indicating strict matching. Plots labelled (av) concern matching to *available* resources; solid lines show mean proportion of time spent on A (ordinate) as a function of F_A (abscissa), standard deviations are shown. Plots marked (ob) concern matching to *obtained* resources; mean proportion of time spent on A (ordinate) is scatter-plotted as a function of proportion of resources obtained from A (abscissa), with best-fit lines superimposed. The equation of each best-fit line is given together with a measure of goodness-of-fit (this ‘error’ measure specifies the range around any point on the line that contains at least 50% of the predictions).

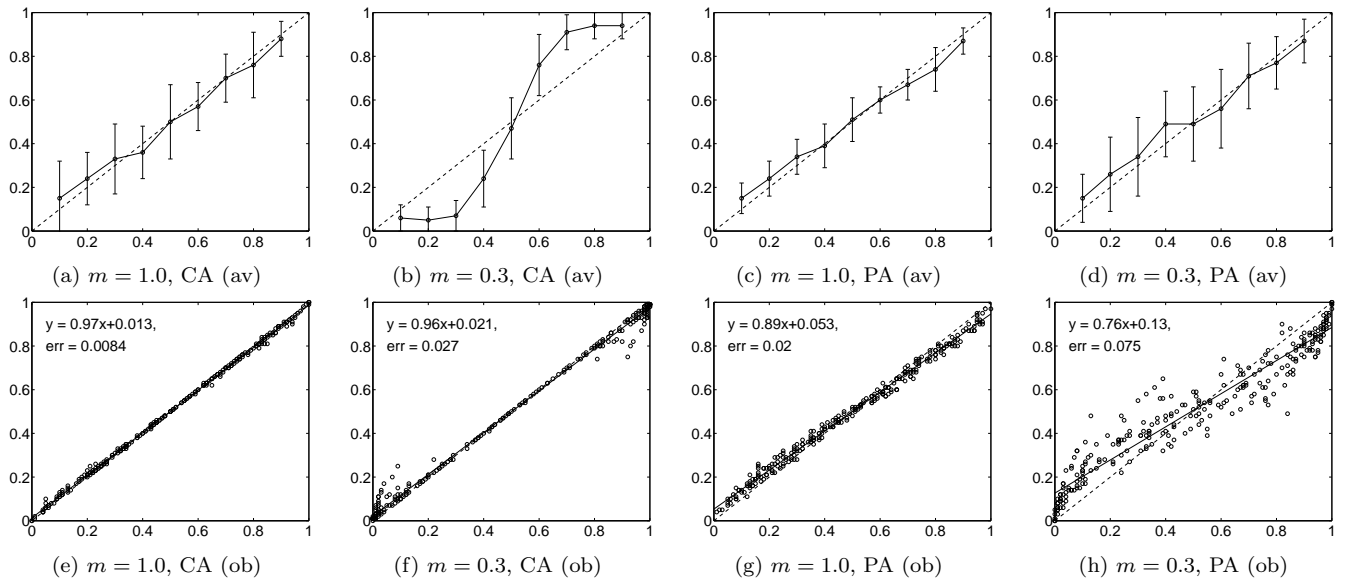


Figure 4: Matching behaviour of embedded ϵ -sampling agents. Dashed lines represent strict matching. Each plot shows mean proportion of time spent on A (ordinate) as a function of resources available (av) or obtained (ob) from A (abscissa), with data collected from 30 analyses at each of 9 values of F_A . Best fit lines are superimposed and equations (with goodness-of-fit) are given as in figure 3.

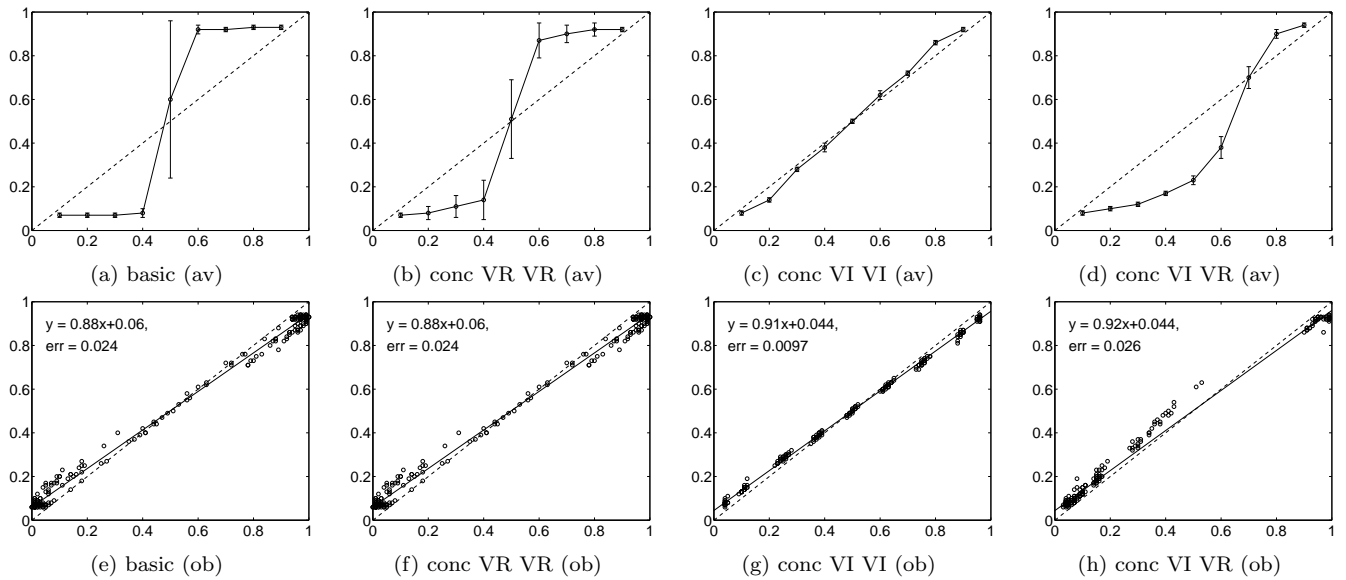


Figure 5: Matching behaviour of isolated ω -sampling agents, to be interpreted as figure 3.




Figure 6: Matching behaviour of embedded ω -sampling agents, to be interpreted as figure 4.

resources, it *cannot* be concluded that isolated agents using the same strategy would fail to match to obtained resources. In other words, matching behaviour in the field may not be a reliable indicator of the performance of isolated individuals with regard to the individual matching law.

As a candidate mechanism underlying both the IFD and individual matching, then, ω -sampling is clearly more successful than ϵ -sampling. This is perhaps unsurprising; by maintaining multiple estimates and employing a more flexible switching rule, ω -sampling is much less likely than ϵ -sampling to be adversely affected by the indeterminacy of P-allocation (with respect to the IFD) or of the conc VR VR, conc VI VI, and conc VI VR schedules (with respect to the matching law), and it is of course in these very conditions that the inadequacies of ϵ -sampling are revealed. One may nonetheless conclude that there do exist simple heuristics capable of underlying both the IFD and individual matching, in several non-trivial situations: ω -sampling is such a strategy, ϵ -sampling is not.

A significant cost of ω -sampling, however, is that it requires concurrent maintenance of more than one value estimate, a cost which would seem rise with environment complexity. Certainly, many ‘mechanisms of matching’ proposed in the psychology literature use multiple estimates (for example, ‘melioration’, Herrnstein, 1982, or ‘momentary maximisation’, Hinson & Staddon, 1983), but is it necessary to assume that agents are able to maintain estimates for *every* patch in their environment? A possible answer is perhaps only to the extent that agents can specify which patch (or option) to choose, which may seem trite, but in fact many patch-switching strategies specify only when to *depart* from a current

patch (or option), without specifying where to go afterwards (Bernstein et al., 1991, for example). Consider the IFD. If only random movement is possible, and all areas of the environment can be accessed with equal ease, then it would only be necessary to maintain a single estimate (to prompt departure). If directed (non-random) movement is possible to any part of the environment (again with equal ease) then concurrent estimates of every patch might be valuable. However, if it is assumed that movement is somewhat restricted, but non-random, then some intermediate solution is likely to be best, at which the agent only maintains a few functionally relevant estimates.

6 Discussion and summary

Perhaps the single most important result of those presented above is that ω -sampling entails matching in at least one situation in which such behaviour is suboptimal (conc VI VR). This supports the idea that some instances of individual suboptimal behaviour can be understood in terms of the operation of mechanisms adapted to a group context; ω -sampling can be considered *ecologically rational*. Of course, this does not imply that *all* irrational behaviour must be explained this way, and nor are alternative explanations of matching necessarily excluded (Niv, Joel, Meilijson, & Ruppin, 2001, for example), we have simply demonstrated that the present account is at least plausible, and perhaps parsimonious.

In a previous paper this same idea was explored in the context of foraging within a single patch containing both rich and poor resources (Seth, 1999, see also Seth, in press), and a similar result was uncovered: agents evolved in a group context displayed matching behaviour

when assessed in isolation, but those evolved in isolation did not. Compared to the present model, this study was minimal in some ways - no distinct patches, no complex reinforcement schedules, no learning or lifetime adaptation - but rich in others, particularly in the extent to which it captured sensorimotor interactions between agents and their immediate environment, a level of description the present model abstracts away from. Because of these differences, the character of the matching behaviour observed in the two models is also different. In the single patch case, matching was accounted for purely in terms of sensorimotor interaction patterns, to do with inter-agent interference, that acted as historical constraints (Di Paolo, 2001), and was identified only qualitatively. Here, matching has been accounted for in terms of the IFD, and has been identified very closely with its description in the literature. Taken together, these studies light the way to a range of matching phenomena, some of which attach to distinct, separable, and spatially arbitrary choices, others to single patches and basic sensorimotor interaction patterns, and still others to intermediate levels of description, in which sensorimotor interactions may engender patch-switching in the generation of choice behaviour. It remains to be seen what further insights into matching behaviour, and decision making in general, can be attained by modelling at these levels.

A final comment concerns the familiar problem of ‘action selection’, as it is understood in the adaptive behaviour literature (Tyrrell, 1993). An obvious difference between action selection and matching is that the former normally analyses choices that satisfy *distinct* requirements (for example feeding and sleeping), and the latter concerns different ways of satisfying the *same* requirement. Yet the concepts are clearly very close, and no satisfying and general account of decision making can afford to ignore either. To risk belabouring a commonly made point, the psychological literature contains a wealth of conceptual and empirical resources that relatively novel methodologies directed towards understanding adaptive behaviour, for example agent-based modelling, would do well to consider.

In summary, this paper has assessed two simple heuristics, ϵ -sampling and ω -sampling, on their ability to underlie both matching behaviour and the IFD in a number of non-trivial conditions. Contrary to the claims of Thuisjman et al. (1995), ϵ -sampling proved inadequate at both; ω -sampling, by contrast, was successful, even to the extent of entailing matching under conditions in which such behaviour is suboptimal. ω -sampling can therefore be considered an ecologically rational candidate for a ‘mechanism of matching’, and from this vantage matching behaviour - and its potential suboptimality - can be interpreted in terms of decision making adapted to a group environment. The extent to which

	1,CA	.3,CA	1,PA	.3,PA
ϵ -samp., ϵ	.038	.037	.039	.038
ϵ -samp., γ	.071	.069	.038	.030
ω -samp., ϵ	.061	.040	.105	.040
ω -samp., γ	.021	.065	.022	0.030

Table 1: Evolved parameters. Columns labelled by interference level m and allocation method.

other forms of irrational behaviour can be understood in this way, or in similar ways, remains an interesting and open question. Finally, a novel method of parameter setting has been described which may find useful application in other models of this kind.

Acknowledgements

I am grateful to the CCNR, the Neurosciences Research Foundation, and to my anonymous reviewers, for a combination of financial support and constructive discussion. The material in this paper is drawn in part from my D.Phil. thesis (Seth, 2000).

Appendix A

For both ϵ - and ω -sampling, initially random populations (size 100) were evolved in each of 4 cases ($m = 1.0$ or 0.3 , C-allocation ‘CA’ or P-allocation ‘PA’): 8 conditions in total. Each agent (in each condition) possessed a genome of 2 real numbers (range $[0.0, 1.0]$) specifying ϵ and γ . Each condition applied a tournament GA for 100 generations (mutation rate 0.01, each mutation drawn from Gaussian distribution radius 0.13; range transgressions were truncated). Fitness was averaged over 10 separate evaluations. Each evaluation randomly assigned values for F_A and F_B ($F_A + F_B = 1.0$, total resource $F^* = 200.0$ in all conditions), and randomly allocated agents between A and B . Fitness of each agent was determined by total accumulated resources after 1000 cycles. The entire GA process was repeated 10 times in each condition, from which average (condition-specific) near-optimal parameter values were recovered (table 1). Analysis of each condition used these values. Note that analysis of isolated individual matching utilised the average near-optimal parameters across *all* conditions ($\bar{\epsilon} = 0.052, \bar{\gamma} = 0.427$), since no corresponding populations were evolved in these conditions.

References

- Baum, W. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, 22, 231–242.
- Baum, W., & Kraft, J. (1998). Group choice: Competition, travel, and the ideal free distribution. *Journal of the Experimental Analysis of Behavior*, 69(3), 227–245.
- Bernstein, C., Kacelnik, A., & Krebs, J. (1991). Individual decisions and the distribution of predators in a patchy environment II: The influence of travel costs and structure of the environment. *Journal of Animal Ecology*, 60, 205–225.
- Bullock, S., & Todd, P. (1999). Made to measure: Ecological rationality in structured environments. *Minds and Machines*, 9(4), 497–541.
- Charnov, E. (1976). Optimal foraging: The marginal value theorem. *Theoretical Population Biology*, 9, 129–136.

- Conger, R., & Killeen, P. (1974). Use of concurrent operants in small group research. *Pacific Sociological Review*, 17, 399–416.
- Dallery, J., & Baum, W. (1991). The functional equivalence of operant behavior and foraging. *Animal Learning and Behavior*, 19(2), 146–152.
- Davison, M., & McCarthy, D. (1988). *The matching law*. Erlbaum, Hillsdale, NJ.
- Di Paolo, E. (2001). Artificial life and historical processes. In Kelemen, J., & Sosik, P. (Eds.), *Proceedings of the Sixth European Conference on Artificial Life*, pp. 649–658. Springer-Verlag.
- Fagen, R. (1987). A generalized habitat matching law. *Evolutionary Ecology*, 1, 5–10.
- Fretwell, S. (1972). *Populations in seasonal environments*. Princeton University Press, Princeton.
- Gallistel, C., Brown, A., Carey, S., Gelman, R., & Keil, F. (1991). Lessons from animal learning for the study of cognitive development. In Carey, S., & Gelman, R. (Eds.), *The epigenesis of mind*, pp. 3–37. Lawrence Erlbaum, Hillsdale, NJ.
- Goss-Custard, J. (1980). Competition for food and interference amongst waders. *Ardea*, 68, 31–52.
- Gray, R. (1994). Sparrows, matching, and the ideal free distribution: Can biological and psychological approaches be synthesised?. *Animal Behaviour*, 48, 411–423.
- Herrnstein, R. (1961). Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4, 267–272.
- Herrnstein, R. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13(2), 243–266.
- Herrnstein, R. (1982). Melioration as behavioural dynamism. In Commons, M., Herrnstein, R., & Rachlin, H. (Eds.), *Quantitative analyses of behavior, vol II: Matching and maximizing accounts*, pp. 433–458. Ballinger Publishing Co., Cambridge, MA.
- Herrnstein, R. (1997). *The matching law*. Harvard University Press, Cambridge, MA. A posthumous collection of the papers of R.J. Herrnstein, edited by H. Rachlin and D.I. Laibson.
- Herrnstein, R., & Heyman, G. (1979). Is matching compatible with reinforcement maximization on concurrent variable interval, variable ratio?. *Journal of the Experimental Analysis of Behavior*, 31, 209–223.
- Herrnstein, R., & Vaughan, Jr., W. (1980). Melioration and behavioral allocation. In Staddon, J. (Ed.), *Limits to action: The allocation of individual behavior*. Academic Press, New York.
- Hinson, J., & Staddon, J. (1983). Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior*, 39, 25–47.
- Houston, A. (1986). The matching law applies to wagtails' foraging in the wild. *Journal of the Experimental Analysis of Behavior*, 45, 15–18.
- Krebs, J., & Kacelnik, A. (1991). Decision making. In Krebs, J., & Davies, N. (Eds.), *Behavioural ecology*, pp. 105–137. Blackwell Scientific Publishers, Oxford. 3rd edition.
- Milinski, M., & Parker, G. (1991). Competition for resources. In Krebs, J., & Davies, N. (Eds.), *Behavioural ecology*, pp. 137–168. Blackwell Scientific Publishers, Oxford. 3rd edition.
- Myers, J. (1976). Probability learning and sequence learning. In Estes, W. (Ed.), *Handbook of learning and cognitive processes*, Vol. 1, pp. 171–205. Erlbaum, Hillsdale, NJ.
- Niv, Y., Joel, D., Meilijson, I., & Ruppin, E. (2001). Evolution of reinforcement learning in uncertain environments: emergence of risk aversion and matching. In Kelemen, J., & Sosik, P. (Eds.), *Proceedings of the Sixth European Conference on Artificial Life*, pp. 252–261. Springer-Verlag.
- Seth, A. (1999). Evolving behavioural choice: An investigation of Herrnstein's matching law. In Floreano, D., Nicoud, J.-D., & Mondada, F. (Eds.), *Proceedings of the Fifth European Conference on Artificial Life*, pp. 225–236. Springer-Verlag.
- Seth, A. (2000). *On the relations between behaviour, mechanism, and environment: Explorations in artificial evolution*. Ph.D. thesis, University of Sussex.
- Seth, A. (in press). Modelling group foraging: individual suboptimality, interference, and matching. *Adaptive Behavior*.
- Shettleworth, S. (1988). Foraging as operant behavior and operant behavior as foraging: What have we learned?. In Bower, G. (Ed.), *The psychology of learning and motivation: Advances in research and theory*, Vol. 22, pp. 1–49. Academic Press, New York.
- Sutherland, N., & Mackintosh, N. (1971). *Mechanisms of animal discrimination learning*. Academic Press, New York.
- Sutherland, W. (1983). Aggregation and the 'ideal free' distribution. *Journal of Animal Ecology*, 52, 821–828.
- Thuisjman, F., Peleg, B., Amitai, M., & Shmida, A. (1995). Automata, matching, and foraging behaviour of bees. *Journal of Theoretical Biology*, 175, 305–316.
- Tregenza, T., Parker, G., & Thompson, D. (1996). Interference and the ideal free distribution: Models and tests. *Behavioral Ecology*, 7(4), 379–386.
- Tyrrell, T. (1993). *Computational mechanisms for action selection*. Ph.D. thesis, University of Edinburgh.
- Van der Meer, J., & Ens, B. (1997). Models of interference and their consequences for the spatial distribution of ideal and free predators. *Journal of Animal Ecology*, 66, 846–858.