



Patch leaving in humans: can a generalist adapt its rules to dispersal of items across patches?

JOHN M. C. HUTCHINSON*, ANDREAS WILKE*† & PETER M. TODD*‡

*Max Planck Institute for Human Development, Berlin

†Department of Anthropology, University of California, Los Angeles

‡School of Informatics and Cognitive Science Program, Indiana University

(Received 16 February 2007; initial acceptance 30 March 2007;
final acceptance 11 September 2007; published online 27 December 2007; MS. number: 9270R)

We used a computer game to examine three aspects of patch-leaving decisions in humans: how well do humans perform compared to the optimal policy, can they adjust their behaviour adaptively in response to different distributions of prey across patches and on what cues are their decisions based? Subjects earned money by catching fish when they briefly appeared within a pond; the timing of appearances was stochastic but at a rate proportional to how many fish remained. Caught fish were not replaced and ponds varied in how many fish they initially contained (according to three different distributions). At any point subjects could move to a new pond, but travel took some time. They delayed this switch much too long. Furthermore, regardless of the distribution of prey, subjects spent longer at ponds where they had found more items (contrary to optimality predictions in two of the environments). However, they apparently responded not to the number of captures directly (despite this appearing on screen) but to the current interval without a capture, to the interval preceding the last capture, and to the time spent at the current pond. Self-reports supported this order of cue importance. Subjects often left directly after a capture, perhaps an example of the Concorde fallacy. High success rate in the preceding patch decreased residence time and subjects appeared to be learning to leave earlier over the latter two thirds of the experiment. Minimization of delay to the next capture alone might explain some of the suboptimal behaviour observed.

2007 The Association for the Study of Animal Behaviour. Published by Elsevier Ltd.

Keywords: Concorde fallacy; delay minimization; evolutionary psychology; giving-up time; incremental rule; marginal value theorem; optimal foraging; rule of thumb; simple heuristic; task switching

Suppose, like one of us, that you sometimes spend a week turning over thousands of leaves on the forest floor searching for slugs. Twenty minutes can often pass without finding one and before then you would probably already have moved on and tried another likely looking hollow. But if you have found one or two specimens in a particular spot, should you stick there because it appears a good spot or move because it is liable to be exhausted? These conflicting arguments must often have concerned our hunter–gatherer ancestors. The answer depends on how evenly prey are

dispersed across suitable patches. One can gauge this only from foraging success at nearby spots searched earlier, although prior experience elsewhere with similar prey might help suggest which patch-leaving rules work best.

Optimal foraging theory has provided the optimal solution to the problem, although only once the prey distribution has already been learnt. The classic result is the marginal value theorem (Charnov 1976), that foragers should leave a patch when the current rate of return falls below the mean rate under the optimal strategy. However, this rule becomes only an approximation to the solution when patches contain few discrete items, because stochasticity disrupts estimation of current rate of return. Presently we detail the optimal strategies in this case.

Numerous studies have tested how well animals approximate these optimal policies (Nonacs 2001), mostly in ovipositing insect parasitoids (van Alphen et al. 2003; Wajnberg 2006) but also in fish, birds and mammals

Correspondence and present address: J. M. C. Hutchinson, Privat, Staatliches Museum für Naturkunde Görlitz, PF 300154, 02806 Görlitz, Germany (email: majmch@googlemail.com). A. Wilke is at Department of Anthropology, University of California, 341 Haines Hall, Box 951553, Los Angeles, CA 90095-1553, U.S.A. P. M. Todd is at Cognitive Science, Indiana University, 1101 E. 10th Street, Bloomington, IN 47405, U.S.A.

(e.g. Ydenberg 1984; Redhead & Tyler 1988; Marschall et al. 1989; Kamil & Clements 1990; Kramer & Weary 1991; Agetsuma 1999; van Gils et al. 2003). However, we have found scarcely any tests on humans. In this study we used a computer game, based on fishing at a succession of ponds, to measure how well humans perform compared with the optimal solution. We also investigated the cues and rules of thumb used to decide when to switch patches. Thirdly we tested whether humans adjust their strategy to the distribution of prey, as they should. Humans are a particularly appropriate animal for this test because we are extreme generalists and should surely be adapted to forage on some foods that are evenly distributed across suitable patches and some that are concentrated in a minority. Alternatively, rather than facultatively adjusting, do we gamble on the commonest sort of distribution (probably aggregated) or perhaps use a robust rule that performs reasonably well with all distributions?

Previous Work on Humans

Anthropologists have invoked the marginal value theorem to elucidate a variety of decisions of native peoples: when to leave one hunting or fishing ground, when to move herds to fresh pasture, when to cultivate a new patch of forest (Smith & Wishnie 2000), how far food is processed before carrying it home (Metcalf & Barlow 1992) and how much a carcass is butchered (Burger et al. 2005). It is also often invoked as a more parsimonious explanation of cases of apparent conservation, because it predicts switching patches before complete depletion (Smith 1983; Smith & Wishnie 2000). However, anthropologists' tests of whether humans really follow the optimality predictions are qualitative and weak (among the stronger examples are Smith 1991 and Aswani 1998).

The patch-leaving analogy has also been invoked by cognitive scientists to model foraging for information (Pirolli & Card 1999; Pirolli 2005). For instance Pirolli (2005) studied when subjects started anew with an Internet search (cf. switching patches) instead of following links within the current web page (cf. foraging within a patch). However, the model that Pirolli (2005) applied to explain these data was unconnected with those considered in behavioural ecology.

The data most relevant to patch-leaving decisions of humans come from operant experiments (Wanchisen et al. 1992; Hackenberg & Axtell 1993). Subjects had to press a key numerous times to get a reward. For one option (progressive ratio) the required number increased with each trial (cf. patch depletion). For the alternative option (fixed ratio) the required number was fixed, and, crucially, choosing this option reset the number of presses required by the progressive-ratio option to the minimum (cf. travel to a fresh patch). Subjects responded adaptively to changes in the fixed-ratio requirement and often found the optimal response. Results were similar when subjects had merely to wait particular times after selecting an option before receiving each reward (progressive interval; Hackenberg & Axtell 1993).

Evidence of Adjustment to Dispersion across Patches

Patch-leaving decisions are known to adjust facultatively to a variety of aspects of the environment (e.g. previous patch quality, prior experience of competitors, perceived mortality risk, travel time: Roitberg et al. 1993; Visser et al. 1992; Cuthill et al. 1994; Goubault et al. 2005; Tentelier et al. 2006; Wajnberg et al. 2006). Also, different species and populations of parasitoid wasps are hardwired to use different patch-leaving rules that would adapt them to different degrees of host dispersal across patches (Wajnberg et al. 1999, 2003). But few studies have tested for a similar facultative adjustment.

Vos et al. (1998) studied switching behaviour in a parasitoid of two congeneric hosts, one with an aggregated dispersion and one more solitary. The cues that were significant in each environment (host) differed appropriately (although the difference was not formally tested). However, the parasitoids might feasibly use host identity as a cue to switch between alternative hardwired rules, whereas our human subjects could learn the distribution and appropriate rules only from their own foraging success on a novel resource. Three experiments involving parasitoids, fish and birds have failed to detect such an ability (Marschall et al. 1989; Dall 1995, chapter 5; Burger et al. 2006). But, although one experiment in blue jays, *Cyanocitta cristata*, also failed to find a difference (Kamil & Yoerg 1985), another suggested greater dependence on time since the last capture when patch quality was more variable (Kamil & Clements 1990). Two other examples also show an ability to change the form of a rule in a way which would enable adaptation to different dispersions (see Theory and Predictions), although other factors triggered the change. Great tits, *Parus major*, switched either after a fixed number of rewards or after a particular giving-up time, depending on whether intercapture intervals were deterministic or stochastic (Ydenberg 1984). And in one insect parasitoid species ovipositions either increased or decreased tendency to switch, depending on egg load (Outreman et al. 2005).

THEORY AND PREDICTIONS

Our experiment involves a solitary forager encountering prey randomly in a patch, with encounter rate proportional to number of prey remaining. Iwasa et al. (1981) showed that, if the forager knows the distribution of the number of items per patch, the a posteriori probability that n items remain depends on only the number of items caught in the current patch (N) and the time spent on it (T). The optimal policy when to switch patches thus also depends on only N and T . There is no additional advantage to remembering time since the last prey capture (I) or timings of earlier captures. Given this, the optimal policy may be calculated by dynamic programming (e.g. Green 1980, 2006; Olsson & Holmgren 1998).

The optimal policy specifies whether to leave or stay given each combination of N and T and its form depends on the distribution of items per patch (Iwasa et al. 1981;

Green 1987). If the number of items per patch is Poisson distributed, the forager should leave each patch after a fixed time (excluding times spent handling prey) independent of foraging success ('fixed- T rule'). With more aggregated distributions than Poisson, each capture should extend residence time, because a capture suggests high patch quality ('incremental rule'). Conversely with distributions more dispersed than Poisson, each capture should decrease residence time ('decremental rule'). The extreme case is an even distribution of items across patches, where it is optimal to leave once a fixed number of items (not in general all) are caught, regardless of when this happens

('fixed- N rule'). Figure 1a shows the optimal policies for the four environments used in our experiment (detailed under Methods).

The fixed- N and fixed- T policies appear particularly easy to follow, and the other cases may be approximated by simple mechanisms analogous to a clockwork egg timer (Waage 1979; Ollason 1980; Green 1984; Stephens & Krebs 1986, chapter 8; Pierre et al. 2003). When the patch is entered the timer is wound up a certain amount; it then runs down until the 'ping' triggers leaving, but meanwhile each capture leads to extra winding up (incremental rule) or down (decremental).

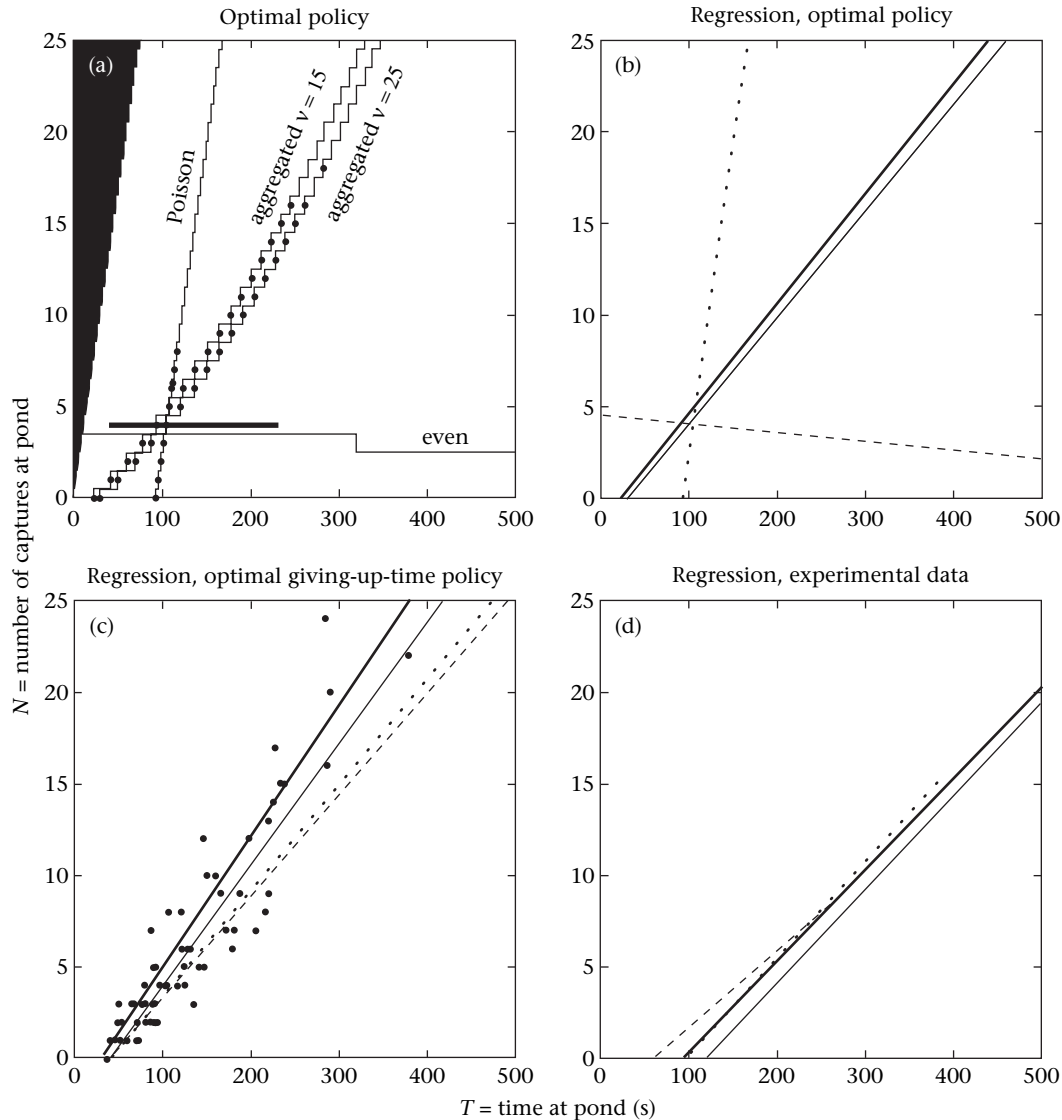


Figure 1. (a) Optimal switching lines for each environment: on arriving at a pond the position in N - T state space is at the origin and switching should occur when the appropriate switching line is crossed. Dots and the thick black line show the range of switching points exhibited in 95% of 5×10^5 simulated patch visits. The line for the even environment would be horizontal were there not slight variation in patch quality associated with the mean quality not being integral, and the line for the Poisson environment would be vertical were it not for the 3-s handling time (see Methods). Handling times make the black region to the left inaccessible. (b) Regression lines fitted to the switching points generated by the simulations of the optimal policy. Dashed line = even environment; dotted line = Poisson; thin solid line = aggregated with travel time $v = 25$ s; thick solid line = aggregated with $v = 15$ s. (c) As (b) except that the policy is the giving-up-time rule optimal for each environment. The dots are 100 giving-up points generated by simulations of the giving-up-time rule optimal in the aggregated $v = 25$ environment. (d) As (b), except that the switching points were those observed, and the regression includes subject as a random factor.

We have also examined whether subjects used other decision heuristics, possibly based on cues besides N and T . A prominent example is to switch after a fixed period of time without a capture, or since entering the patch if that is later (Iwasa et al. 1981; Green 1984, 1987); this is the 'giving-up-time rule' (McNair 1982). In the environments in our experiment, this rule can perform 88–97% as well as the optimal policy (Table 1). Its best performance was in the aggregated environments, because the giving-up-time rule resembles an incremental rule (optimal there) in that each capture postpones departure. Thus a giving-up-time rule also generates a positive relationship between N_S and T_S (the values of N and T at switching); the scatter of points generated by the best giving-up time has a slope similar to that of the optimal policy in our aggregated environments (Fig. 1c, Table 1).

Another reason to consider alternative cues, even though they are in theory unnecessary, is that the simple policies in Fig. 1a are optimal only when the distribution of patch quality is known. Initially, and probably throughout our short experiment, subjects should try both to make good decisions with imperfect knowledge of this distribution and to learn about it (or try out different rules) so as to perform better in the future. In this situation experience at earlier patches should influence the strategy subsequently applied. Information on intervals between captures is still superfluous if N and T on the current and past patches are properly taken into account. Nevertheless the optimal policy is not so simple as that in a known environment, so simpler rules of thumb using alternative cues may provide more attainable solutions.

In this context of uncertainty about the environment, the giving-up-time rule sometimes has a strength in being robust. Table 2 shows the performance of the policy optimal in our even environment when transferred to the aggregated environment and vice versa. Performances are worse than those for a giving-up-time rule or fixed- T rule adapted to one environment but transferred to the other. The giving-up-time rule is not always more robust (e.g. when travel times are swapped: Table 2; Green 1984 provides another example). However, because giving-up-time rules measure rate over a shorter time window than N and T , they should be more robust to within-patch environmental changes, as when suddenly all prey hide in a refuge (cf. Shettleworth & Plowright 1992).

METHODS

The 100 human subjects (64 women, 36 men; age ranges 19–30 and 22–31 years, means 25 and 27 years) carried out the task in a laboratory dedicated to such psychological experiments. The complete experiment, including instructions and debriefing, was presented on a computer (Czienskowski 2005). The task was to catch as many fish as possible by the end of the experiment; catching the fish that occasionally appeared was easy but the skill came in judging when to switch from a depleted pond to a fresh pond. Payment was exclusively contingent on performance but delivered only at the end.

The main screen showed a cartoon character fishing at a pond, with the float on the end of the fishing line forming the mouse cursor (figured in the Supplementary Material). Fish popped to the surface stochastically at a rate of $0.005 \times (\text{number of fish in pond}) \text{ s}^{-1}$. Thus the rate decreased as the subject depleted a pond's stock (i.e. random, rather than systematic, search: Green 1987; Marschall et al. 1989). Fish were somewhat cryptic and remained visible for only 2 s, so subjects had to remain attentive to catch fish, by clicking on them with the mouse. If a fish was caught, it moved to the bucket (an animation lasting 2 s). During this transfer no further fish would appear and the mouse was inactivated (so effectively there is a handling time of 2 s plus the 1 s typically taken to click on a fish). Each capture also added a fish icon to a panel which thus displayed the number of fish caught at the current pond. Subjects had to rely on their own perception of time (watches and telephones were not allowed).

Subjects could switch ponds at any time by clicking on a prominent labelled panel. This would cause the angler to walk off screen, followed by a delay in which a bouncing-ball animation with a 'Please wait!' sign was shown. Then a new pond appeared and the angler walked up to it and stuck out the fishing rod. To encourage the perception that it was a fresh pond, at each switch we rotated the pond border and placed three tussocks of vegetation randomly around it.

We used a between-subjects experimental design with subjects randomly assigned in equal numbers to four environments. Environments differed only in the distribution of fish across ponds or, in one case, in the travel time between ponds. The mean of each distribution was 9.7 fish.

Table 1. Performance under the optimal and constrained-optimal strategies

Environment	Maximum gain rate (fish/s)	Giving-up-time rule		Fixed- T rule	
		Optimal threshold (s)	Relative performance	Optimal threshold (s)	Relative performance
Even	0.0293	40	0.88	93	0.95
Poisson	0.0280	39	0.96	93	1.00
Aggregated $v=25$ s	0.0312	36	0.97	93	0.90
Aggregated $v=15$ s	0.0346	31	0.97	74	0.89

The optimal policy is compared against rules in which subjects switch after a fixed giving-up time without a capture or after a fixed time in the patch (excluding handling times). These constrained optima were calculated by simulation (5×10^6 patches) to an accuracy of 0.1 s; their relative performances are given as a fraction of the long-term gain rate under the optimal policy.

Table 2. Robustness of optimal and two constrained-optimal strategies

Environment to which adapted	Environment in which assessed	Performance relative to maximum		
		Optimal	Giving-up-time	Fixed-time
Even	Aggregated $\nu=25$ s	0.41	0.97	0.90
Aggregated $\nu=25$ s	Even	0.85	0.87	0.95
Aggregated $\nu=15$ s	Aggregated $\nu=25$ s	0.99	0.96	0.88
Aggregated $\nu=25$ s	Aggregated $\nu=15$ s	0.99	0.96	0.89

Strategy parameters were adapted to one environment and performance was assessed in another. Performances are expressed as a fraction of the long-term reward rate in the new environment under the policy optimal in it. Each value derives from simulations involving 5×10^6 patches.

One distribution was almost even (9 or 10 fish per pond), in another environment fish were distributed randomly over ponds (i.e. Poisson) and in the two other environments the distribution was aggregated following a negative-binomial distribution (variance 4.4 times that for the Poisson distribution, thus more very rich and very poor ponds, but still unimodal). In one of the aggregated environments the travel time ν between ponds (including 2-s animations of the angler moving off and on screen) was 15 s, whereas in the other three environments $\nu = 25$ s. Qualities of successive ponds were uncorrelated (cf. Klaasen et al. 2007).

Subjects played this game in a continuous session which terminated at the first switch after 45 min. This was preceded by a 4-min practice session of identical format except that we informed subjects that it was to allow familiarization with the task, that it was unpaid, so they could feel free to experiment, and that they should try out switching ponds.

Before this practice session, subjects had worked through a series of screens explaining the layout of the main experimental screen, how to catch fish, how to switch ponds and the monetary rewards (0.20 € per fish). The instructions also emphasized two structural aspects of the task. The first concerned the rate of appearance of fish within a pond. Subjects were told that this rate depended only on the number of fish currently in the pond but that the times of appearance were stochastic; only one fish would appear at a time. They were further told that the initial number of fish in each pond could vary and that each capture decreased the number left accordingly (no replenishment).

The second structural aspect concerned switching ponds. Subjects were told that the number of new ponds was limitless but they could not return to ponds visited previously. Their one key task, in addition to catching fish, was to decide when to switch ponds. The only clues about strategy were in the following text (translated from the German). 'Because your final payout depends on the total number of fish caught, you will achieve the best result (= higher payout) if you avoid two extreme strategies: staying too long at each pond, or switching too soon. For instance, switching too soon does not lead to a high performance because of the time delays that follow clicking the switch-pond button.'

After the main experiment, an on-screen questionnaire asked subjects to describe any strategy that they had used to determine when to switch. The next screen asked about their use of the three cues suggested by theory: the number

of fish captured at that pond, the total time spent there and the wait since the last capture. Use of each cue was scored on a four-point scale from 'not at all' to 'very often'.

RESULTS

Testing for Inattention

One concern was that 45 min of the experimental session would generate boredom or tiredness. A potential sign of boredom was mouse clicking in the absence of fish. Four individuals made over 100 such clicks, but they behaved so only in a minority of ponds and were not particularly poor performers. The habit did not consistently increase with time.

Inaccurate clicks or slow reactions might also indicate inattention. One subject failed to catch 16% of fish that appeared, but the median was under 3%. Such failures were associated mostly with not making any click, and in about half of these cases the cursor was not near the fish, so it was probably not spotted. These latter instances occurred a median of only once per subject (1.2% of fish encountered), although three subjects did not react to four or five fish. The median proportion of fish that a subject missed despite a click close by was 1%, and it was always under 6%. The proportion of fish missed, the proportion of fish leading to a close-miss click, and reaction times all tended to fall during the course of the experimental session, but most improvement was in the first 15 min, excluded anyway from most analyses. There was no increase in the rate at which fish were apparently overlooked.

In summary, there is no evidence of a loss of attention over the course of the experiment. If missing the fish was very common, the optimal policy would change quantitatively for two reasons: a lower rate of capture and the forager gaining extra information about the number of fish remaining from having missed one. But because misses are so rare and because ponds are usually left with several fish remaining, behaviour should be scarcely affected. Subsequent analyses ignored misses.

Initial Analyses

In the experimental session a subject encountered a median of 73 fish (range 34–118) from a median of 14

ponds (range 6–31). The median catch at a pond was four fish, with a median residence time of 163 s. The longest residence time was 784 s, resulting in 35 fish caught, the biggest catch from a single pond. Except for a couple of outliers at the start, the shortest residence time was 26 s.

One analysis often considered in the literature is whether more time is spent in a richer patch than in a poorer patch, indicating that foraging success affects the switching decision. This is adaptive in an aggregated environment, but the optimal policy in the Poisson environment is to spend a fixed time (excluding handling times) and the optimal policy in the near-even environment tends to leave richer ponds earlier. We indeed found a positive correlation in the majority of subjects facing aggregated environments, but the number of nonpositive correlations is significantly more than that expected under the optimal policy (Table 3). In the near-even and Poisson environments the frequencies of positive correlations are within the wide ranges expected.

Another recommended analysis tests whether, in environments more aggregated than Poisson, the number of items left after switching is higher if more time has been spent at the patch (e.g. Olsson & Holmgren 1998; Olsson & Brown 2006). This is adaptive because the longer time spent, the better the estimate of current prey density, and thus the threshold (minimum T) for remaining can be set higher without increasing the risk of leaving a good patch too early. Most subjects failed to show the expected positive correlation, although in the aggregated environment with $v = 25$ s the number is within the 95% confidence bounds generated by simulating the optimal policy (Table 3). Really the predicted effect is too subtle to detect readily here. In the even and Poisson environments, roughly zero correlation is predicted under the optimal policy, but observed correlations are almost all negative: subjects spent longer times if they had caught more fish, thus depleting such ponds to lower levels.

To elucidate individual policies, we plotted each subject's trajectories in N vs T space to compare against the

optimality predictions in Fig. 1a, b: Fig. 2a shows a random selection of subjects. The predictions hold to the extent that the switch points often lie roughly on a straight line. However, much intersubject variability is apparent in the slope and position of such lines.

Figure 2b plots the same histories in I vs T space, where I is the interval without a capture (or from arrival at the pond, if this is later). A simple giving-up-time rule would place all switches along a horizontal line, which was not crossed earlier. The data are poorly characterized by either prediction.

Also, in all environments the distribution of giving-up times is strongly bimodal with peaks either side of 4 s (Fig. 3): the proportion of switches within 4 s of a capture was 0.20 and 44 subjects showed at least one such switch (these values ignore the first 15 min). However, in only 12 subjects were the majority of switches like this, and in only one subject were all switches so: this subject was in the even environment where such behaviour is predicted, but he did not switch after a constant number of captures. Our intuition was that a long period without a capture might make subjects promise themselves to switch directly after the next capture. Indeed, switches within 4 s of a capture were strongly associated with a long fruitless wait beforehand (the median interval preceding a <4-s switch was 45 s, compared with 21 s for other switches). We consider alternative explanations in the Discussion.

Regardless of whether we include giving-up times under 4 s, in the aggregated environments giving-up times are negatively associated with initial pond quality ($P < 0.0001$; based on regression including subject as a random factor, with the first 15 min excluded). In the Poisson environment the relationship persists but is significant only when including giving-up times under 4 s. The negative relationship agrees with the finding of Roche (1996) in the black-capped chickadee, *Parus atricapillus*; he shows that the result is compatible with a variety of patch-leaving rules. However, the optimal policy generates such a relationship only in our Poisson environment and not in the aggregated environments where we observed it.

To conclude, although I might be used as one cue in a more complex method of deciding when to switch (as later analyses confirm), Fig. 2 suggests that it makes sense to start our analysis by assuming a switching line in N vs T space. Even if N and T are not the cues used to determine switching, it is useful to describe the resultant behaviour in terms of the values of N and T at which switching occurs (N_s and T_s), because these determine the payoff. We have assumed that this switching line is straight, based on inspection of Fig. 2a and because this is at least a good approximation to the optimal policies in Fig. 1a. To fit a straight line we regressed T_s on N_s , thus assuming that all error is in the perception of T (subjects had no access to clocks but could see the stack of caught fish). The T_s vs N_s slope may be readily interpreted as how long each capture extends residence time. Its reciprocal gives the N_s vs T_s slope of the lines shown in Figs 1b and 2a. As expected from Weber's law applied to time estimation (e.g. Brunner et al. 1992), residuals increased proportional to the fitted values of T_s ; accordingly our regression analysis specified this variance structure. This, as well as

Table 3. Number of subjects (of 25 in each environment) showing a positive correlation between time foraging at a pond and (1) initial number of fish in pond or (2) number of fish left after switching

Environment	T_{S-} vs initial number of fish		T_{S-} vs number of fish left	
	Observed	Optimal	Observed	Optimal
Even	14	5–14	0	6–15
Poisson	16	7–18	1	7–18
Aggregated $v=25$ s	18	22–25	11	9–18
Aggregated $v=15$ s	20	22–25	8	12–20

Time spent at a pond T_{S-} here excludes handling times. The ranges labelled Optimal are the 95% confidence limits based on 1000 simulations following the optimal policy given the observed number of ponds visited by each subject and the observed number of fish in each pond (but with new intercapture intervals randomly generated according to the specified N -dependent rate). The fixed- T rule optimal in a Poisson environment implies correlations of 0, which are allocated randomly as positive or negative. Analyses omit ponds visited in the first 15 min.

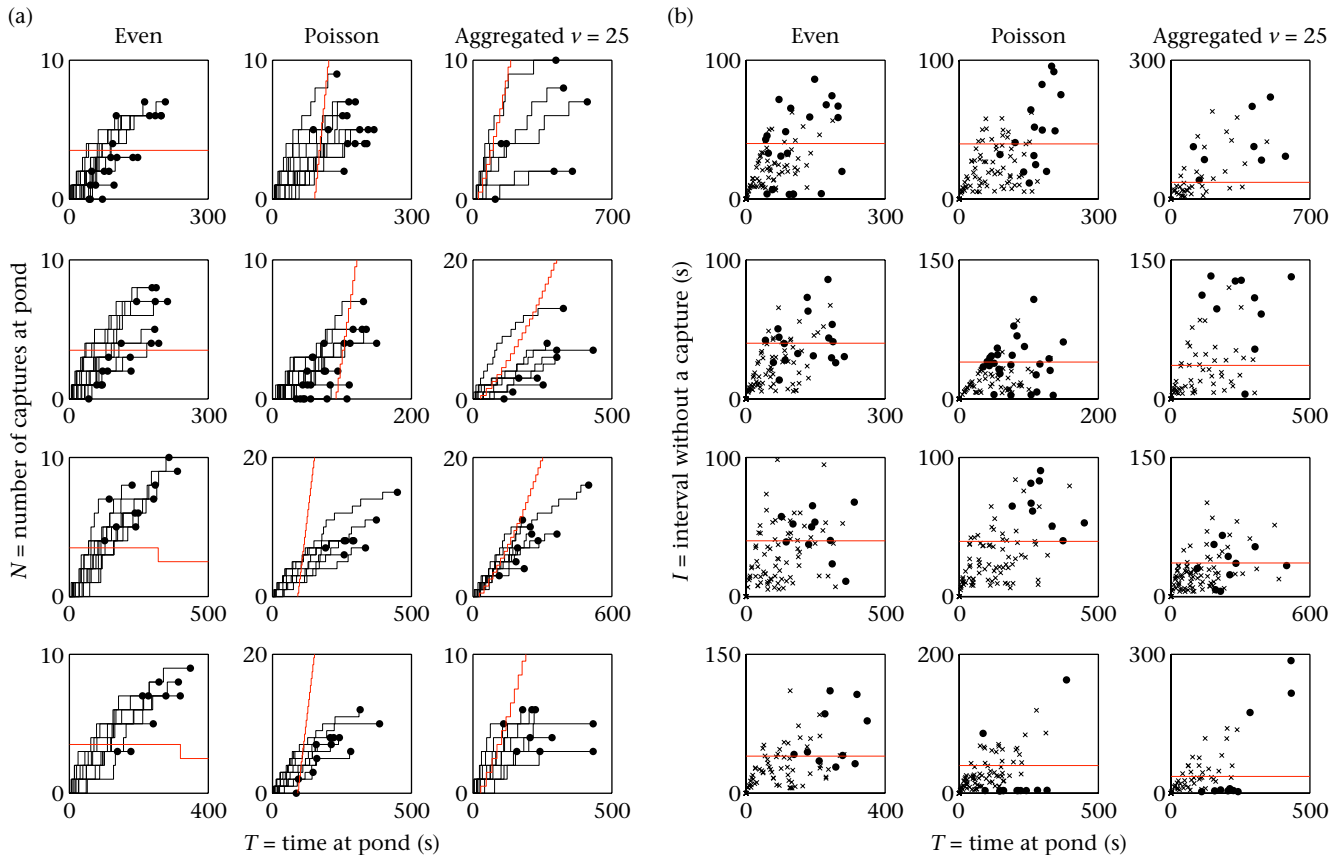


Figure 2. (a) Trajectories in N - T space of 12 randomly chosen subjects from three of the environments. Each staircase maps foraging success at one pond: vertical jumps correspond to prey captures and dots mark the switches. Red lines show the switching line under the policy optimal in each environment (Fig. 1). (b) Duration of intervals without a fish appearing plotted against time at the pond for the same 12 subjects. Crosses mark intervals ending in a capture, dots mark those ending in a switch. Red lines show the giving-up times that maximize performance. Note the variation in scales between subjects; all scales are linear.

treating subject as a random factor in a nested design, is possible using the `lme` function of the statistical program R; it fits using maximum likelihood or restricted maximum likelihood (Pinheiro & Bates 2000).

To describe the position of the switching line, rather than its slope, we considered the average residence time. Because of Weber's law, we log-transformed values of T_S before analysis but report backtransformed values (i.e. our 'average T_S ' is the geometric mean).

Change in Decision Rule through the Experimental Session

Subjects will require experience with an environment to adapt their patch-leaving rules. A mean of only 2.7 ponds were visited during the 4-min practice session, so we expect any adjustment to continue during the experimental session. The direction of any adjustment provides a test of whether subjects adapt to different environments, but we also wanted to facilitate other comparisons between environments by selecting a period when strategies had stabilized.

Our approach was to fit straight-line relationships between T_S and N_S , as outlined above. Each subject was

allowed a different slope and intercept but these were each further affected by an additional adjustment factor, common to all subjects from each environment, that coded for arrival time at each pond (since the start of the experimental session). This factor was categorical with a level for each of six successive 450-s intervals.

Figure 4a shows how the adjustment factor alters the average T_S vs N_S slope over the course of the experimental session. Surprisingly, the pattern is of much initial variation between environments but convergence by the second half of the experiment. If we consider arrival time as a continuous rather than a categorical variable (i.e. fitting straight-line relationships to Fig. 4a), two environments (Poisson and aggregated $v = 25$ s) show significant changes in the T_S vs N_S slope through the course of the experiment ($P = 0.001, 0.017$). The variation between the four environments in the direction and extent of these changes is highly significant ($P < 0.0001$).

Turning to the position of the switching line along the T_S axis, no environment shows a significant change in the average time at a pond during the course of the experimental session ($P > 0.18$), but incorporating all environments into one ANOVA does indicate significant variation between times ($F_{5,1276} = 3.75, P = 0.002$). Figure 4b suggests that this involves an initial increase and then a possible decrease.

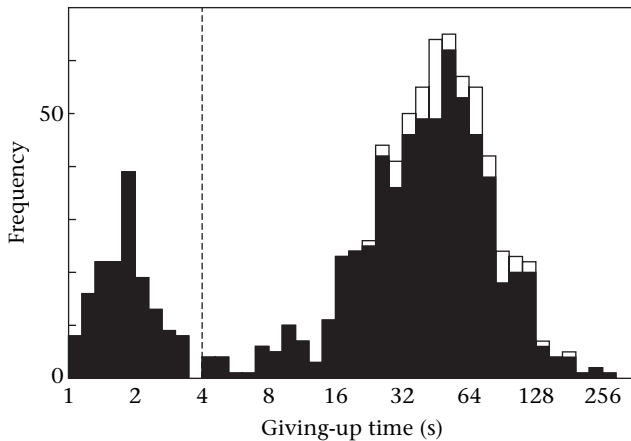


Figure 3. Distribution of giving-up times = time preceding click on 'new pond' button since cursor reappeared at centre of pond either after the last capture (black part of bars) or, if pond was left without any capture, on first arrival at pond (white part of bars). The dashed line at 4 s shows the division point used in the Cox regression (Table 5). Note the logarithmic time scale. Data from ponds visited in the first 15 min are omitted.

Dall (1995, chapter 5) rejected the hypothesis of adaptation to prey dispersion over patches solely on the basis of failing to find an improvement in performance (rewards per probe) over successive experimental sessions. In our experiment, there was also no general improvement in performance (regressing N_S/T_S against arrival time at each pond: $F_{1,1280} = 0.37$, $P = 0.5$) although there was significant temporal variation caused by performance tending to fall and then rise, mirroring the pattern for average T_S .

So as to remove the period of most pronounced strategy experimentation and adjustment, subsequent analyses ignore pond visits starting within 15 min of the start; this leaves a median of eight pond visits for each subject. An exception was made for one subject who would otherwise have been represented by only two ponds; we included a pond at which she arrived after 860 s.

Terminal Relationship between T_S and N_S

Slope of switching line

The environments do not differ significantly in the slope of the T_S vs N_S regression lines ($P = 0.55$, based on $N_S \times$ environment term, with subject as a random factor affecting intercept and slope). The common slope is such that each capture increases residence time by a mean of 19 s (95% confidence interval = 18–20 s). If this value were just 5 s less in the even and Poisson environments than in the aggregated environments, the chance of finding significant ($P < 0.05$) interenvironment variation would be 95%, so statistical power is sufficient to detect differences in slope much smaller than predicted.

Figure 5a shows the distribution of the T_S vs N_S regression slopes for individual subjects. They are strikingly consistent in being positive and mostly incompatible with optimal policies in Poisson as well as even environments. Instead, slopes in all environments group around that predicted as optimal for the particular negative-binomial dispersion that we specified for our aggregated environments.

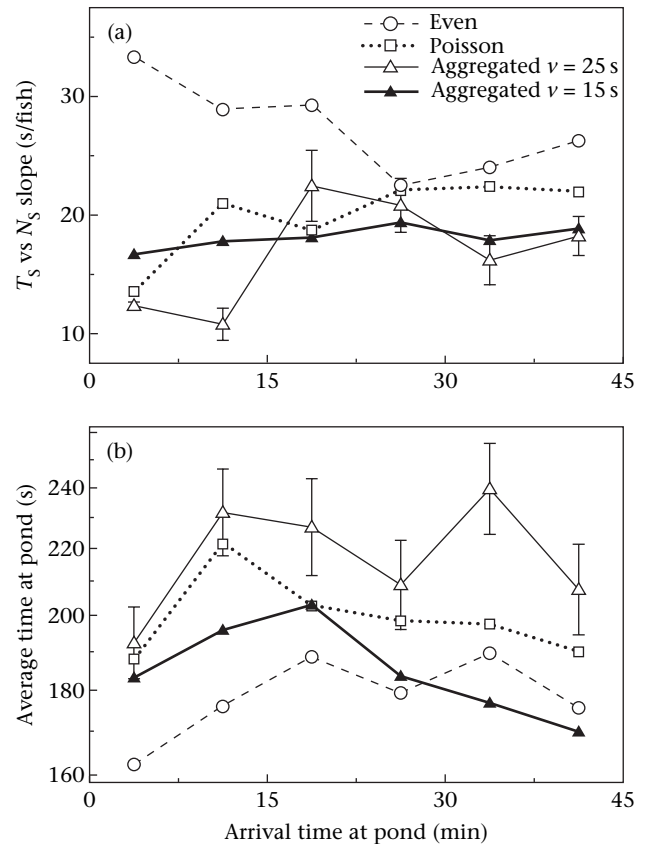


Figure 4. (a) Variation in slope of switching line (increment in residence time with each capture) over the course of the experimental session. Values derive from linear mixed-effect regressions (lme from R) from the coefficients describing the interaction of a six-level factor describing arrival time at a pond (pooling over 450-s intervals) with N_S ; the arrival-time factor also affects the intercept and the model includes subject as a random factor affecting slope and intercept. (b) Variation in the geometric mean of T_S during the experimental session. The values derive from ANOVAs of $\log(T_S)$, in which the arrival-time factor is crossed with a random subject factor. Each environment was analysed separately. The error bars (for one representative environment) show the standard errors appropriate to compare between times within that environment.

We should consider alternative explanations for these positive slopes. If subjects spend random amounts of time at a pond regardless of capture success (cf. Adler & Kotar 1999), the longer time spent the more fish will turn up. To make predictions from this null model, we simulated how many fish would be caught in the observed set of residence times of each subject in ponds of random quality (selected according to their frequency in the environment faced). The resultant data were analysed by a regression procedure identical to that of the real data, and we repeated the simulations 20 times. In three of four environments the observed T_S vs N_S slopes differ significantly from these simulations (Table 4; $P < 0.012$, t tests).

Another explanation for the positive regression slopes is that subjects are using a giving-up-time rule (Fig. 1c). Such a rule optimized to each environment generates slopes of 13.9–18.0 s/capture, which are less than those observed (Table 4). However, the observed slopes can be generated by use of longer giving-up times than optimal (the rule

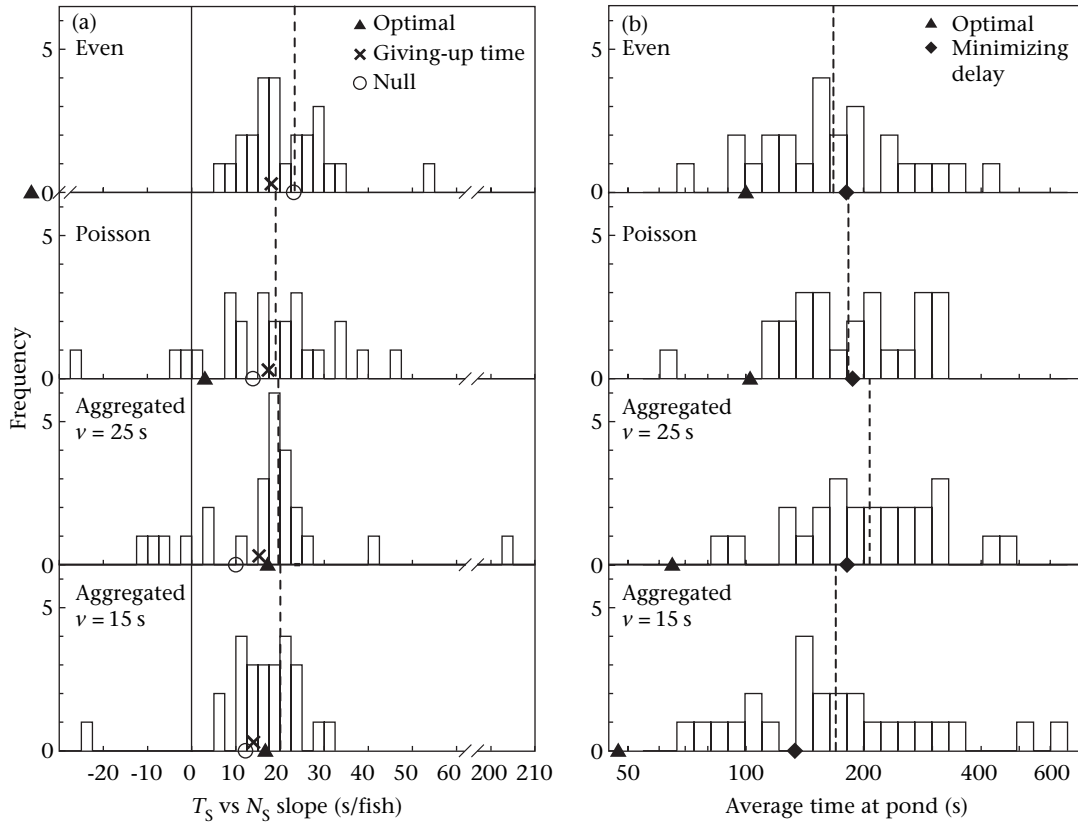


Figure 5. (a) Distribution of regression slopes (increment in residence time with each capture) fitted to data from individual subjects from each of the four environments. Note the break in the axis to accommodate one outlier. The dashed line marks the slope from a regression model combining data from all 25 subjects in an environment, with subject included as a random factor affecting intercept and slope. (b) Distribution of the geometric mean of T_S for each subject. Note the logarithmic scale. The dashed line marks the geometric mean of the subject geometric means. Symbols mark predictions of various models described in the text: note that the optimal slope in the even environment really lies much further off-scale to the left at -211 s/fish. Analyses omit ponds visited in the first 15 min, so for some subjects slope and average T_S are based on only three data points.

would have to be probabilistic to be compatible with the variation in Fig. 2b).

Position of switching line

Figure 5b shows the distribution of average residence time (T_S) for individual subjects. For nearly all subjects (96/100) average T_S is longer than optimal. For instance

in the even environment the average of subjects' average T_S is 168 s compared with 100 s under the optimal policy; in all environments these differences are significant (t tests: all $P < 0.0001$). The observed values are also much higher than those predicted by an optimal giving-up-time rule (Table 4) or by a rule optimal under the constraint that captures delay switching to the extent implied by the observed T_S vs N_S slopes.

Table 4. Testing predictions of optimal and other rules

	T_S vs N_S slope (s/fish)				Average T_S (s)			Relative performance
	Optimal	Optimal if GUT rule	Null model	Observed	Optimal	Optimal if GUT rule	Observed	
Even	-211	18.0	22.7	23.3	100	84	168	0.84
Poisson	3.0	17.3	13.9	19.0	103	83	183	0.86
Aggregated $v=25$ s	17.2	15.2	10.0	19.6	65	72	207	0.73
Aggregated $v=15$ s	16.7	13.9	12.2	20.1	47	58	170	0.72

Slopes are fitted using linear mixed-effect regressions of T_S against N_S . For the observed data, subject is considered a random factor affecting slope and intercept; the slopes given are based on separate analyses of each environment (a joint analysis yields similar results). GUT: giving-up-time. The null model assumes that choice of T_S is independent of N and predictions are based on simulations applying observed values of T_S to ponds of random qualities. Values for average T_S are geometric means (and for observations are the geometric mean of subject geometric means); the optimal values were calculated by simulation. Relative performance is the observed mean rate of gain (N_S and T_S each pooled over subjects) divided by that obtainable under the optimal strategy. Analyses omit ponds visited in the first 15 min.

Environments do not differ significantly in average T_S (ANOVA: $F_{3,96} = 1.13$, $P = 0.3$). However, in the aggregated environments average T_S under the optimal policy is smaller (Table 4). Consequently, performance relative to the maximum achievable is lowest in the aggregated environments (Table 4), despite the positive T_S vs N_S slopes observed being adaptive in that environment. Comparing the two aggregated environments, the optimal policy is to switch on average 18 s earlier when the travel time is 15 s between when it is 25 s. The observed difference in average T_S between the two environments is of the right order (37 s), but considerable variation between subjects prevents it differing significantly from 0.

Table 4 shows observed performances relative to those of the optimal strategies. Note that subjects perform worse than an appropriate giving-up-time or fixed- T rule could achieve, often considerably so (Table 1).

An Alternative Model: Cox Proportional Hazards Regression

We earlier relied on visual inspection of Fig. 2 to choose an analysis in which switching depended on number of captures (N) rather than on interval without a capture (I). A more formal comparison is desirable. The analysis that has become standard is Cox proportional hazards regression (e.g. Haccou et al. 1991; Driessen & Bernstein 1999; Tenhumberg et al. 2001). This assumes that switching is a stochastic process but the rate is affected by various factors. Much of the error in the earlier regression analysis is subsumed into the fitted function h_0 describing the rate of switching in each time unit averaged over all patches. We take this baseline hazard h_0 to be a function of time spent on the pond (some studies instead take I ; such an analysis yielded results compatible with those reported below). Other factors such as N and I can be added to the model and each has a consistent multiplicative effect on the switching tendency at all times: for instance,

$$h(T, N, I) = h_0(T)\omega_i \exp(\beta_N N + \beta_I I).$$

Here h is the hazard function (the probability of switching in one time unit if not switched at the start of the time unit). The fitted coefficients β_N and β_I reveal the effects of N and I on the baseline hazard rate. The coefficient ω_i is a 'frailty' factor describing the random variation between subjects in their tendency to leave (Therneau et al. 2003); ω_i are assumed to be sampled from a gamma distribution. Our analysis used the 'survival' package of the statistical program R (Therneau 1999). Because covariates such as I and N change through the course of a pond visit, a separate line of data was entered to describe state at each second. When a fish appeared but escaped, we excluded the period from the appearance until after the next capture or switch. Statistical significance was assessed by analysis of deviance.

Although this statistical model might well provide a reasonable description of patch-leaving behaviour, it cannot exactly fit either a simple N -dependent rule exemplified by the optimal strategy or a giving-up-time rule. These rules specify that the threshold value of T , applied deterministically, shifts depending on N or I . Cox

regression does not allow factors to shift a threshold in time (unless the factor specifies strata, but then no statistical comparison between levels is possible), but only to multiply the probability of switching by the same factor at all times. Moreover, when we put data simulated from an optimal rule (with added noise) into a Cox regression, I often was significant even though in the model generating the data it had no influence. But more reassuringly, when we entered both N and I together into that analysis, only N came out as significant. And the converse was true when analysing simulated data generated by a noisy giving-up-time rule.

We initially entered the following predictor variables into the Cox regression analysis: N captures from the current pond, I interval without a capture, I_{-1} interval without a capture leading up to the previous capture, I_{-2} intervals without a capture leading up to the second-last capture (only intervals in the current pond count) and t_a arrival time at the pond since the start of the experimental session. We did not include a factor describing true patch quality, because subjects could only estimate this and our aim was to discover what cues they used. Nor did we include total fish caught on all ponds or number of ponds visited, because of the risk that the direction of causation could be reversed (e.g. assuming some autocorrelation in strategy used in successive ponds, the strategy will have affected the number of ponds visited, which confounds any effect of number of ponds visited on the strategy). For the same reason we restricted measures of time and success on the previous pond to a supplementary analysis reported below. We performed a stepwise procedure to remove non-significant variables ($P < 0.05$), at intermediate stages manually checking for the linearity of effects with plots such as those in Fig. 6 and where necessary transforming. We also tried nonlinear combinations of the variables where suggested by these plots or by a priori mechanistic or functional considerations. This explorative flavour increases the risk of type I errors, but all variables in the final model are highly significant (Table 5). Data from all environments were entered into a single analysis, and only after selecting a set of significant predictors did we test for differences between environments.

The fit improved dramatically if we distinguished cases of switching within 4 s of a capture (see Fig. 3). We specified this in the Cox regression by splitting variables such as I_{-1} into two variables I_{-1}^+ and I_{-1}^- , which were both entered into the regression: for a 4-s period following each capture $I_{-1}^+ = 0$ and $I_{-1}^- = I_{-1}$, whereas otherwise $I_{-1}^+ = I_{-1}$ and $I_{-1}^- = 0$. The same technique can cope with the situation before any captures, when I_{-1} is undefined, and before two captures, when I_{-2} is undefined (e.g. when $N = 0$, set $I_{-1} = 0$ and a new factor $X_0 = 1$; otherwise $X_0 = 0$; both I_{-1} and X_0 are included in the Cox regression; the coefficient of X_0 is R_0 in Table 5). We now explain the results summarized in Table 5 and Fig. 6.

The baseline hazard function indicates that time spent on the patch tends to increase the tendency to leave (Fig. 6a). Although the effect is quite weak (a factor of only about four) and is unstable above 400 s, a separate Cox regression with the baseline hazard a function of I confirmed that this effect of T was significant ($P < 0.0001$).

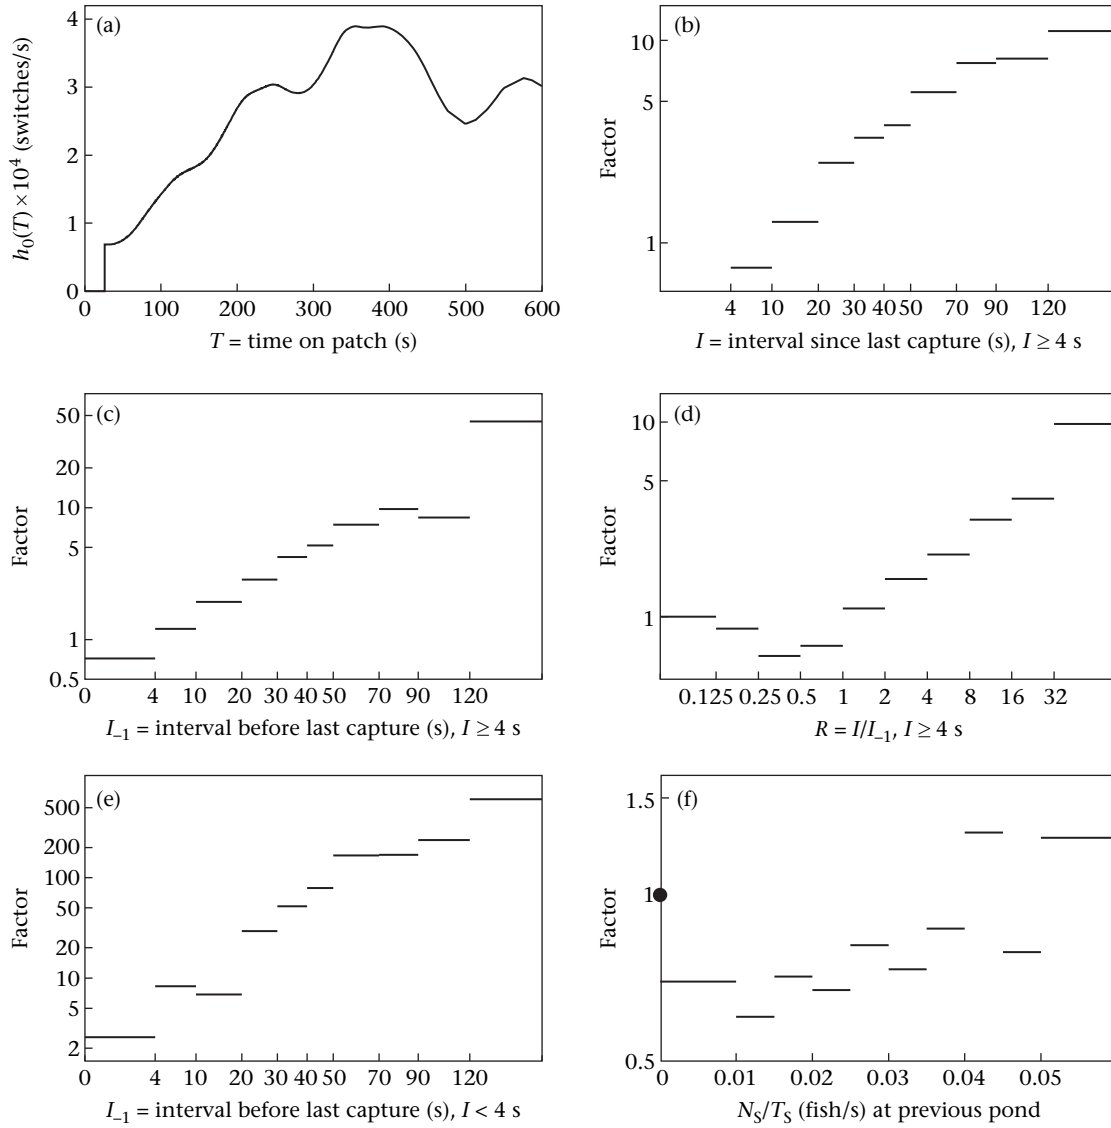


Figure 6. Cox proportional hazards regression. (a) Smoothed version of the baseline hazard function, derived by fitting a 12-parameter spline to the survival function and differentiating. Note that there are relatively few observations above 400 s residence time. (b)–(f) Influence of various other factors on leaving tendency (see also Table 5). In each case we reran the Cox regression with one continuous variable replaced by a counterpart in which it was split into a series of categorical levels; the exponent of the fitted coefficients of each level are plotted on the vertical axes and indicate by what factor this range of values increases leaving tendency over the baseline. Horizontal axes are scaled according to the transformation used in the model in Table 5. In (e) we did not allow the differences between environments present in the full model (Table 5). In (f) the dot on the vertical axis covers cases where the previous pond was left without a capture ($N_s = 0$). The model assumes that a different set of factors influences leaving tendency within 4 s of the last capture compared to that after.

When more than 4 s has passed since a capture, the probability of leaving increases with the interval since the last capture (I). A square-root transformation makes the relationship more linear (Fig. 6b). The interval preceding the last capture (I_{-1}) has an effect similar to that of I (Fig. 6c). In addition there is an effect of the ratio between these two intervals ($R = I/I_{-1}$), so that a decline in capture rate stimulates switching; the effect is nonlinear, with no incremental effect when the current interval is less than the preceding interval (Fig. 6d). However, the form of this relationship depends considerably on the transformations applied to I and I_{-1} , so it is safer to conclude merely that I and I_{-1} both have effects and interact in some

nonadditive way. The interval before the second-last capture (I_{-2}) was not significant.

When less than 4 s has passed since a capture, the interval before that capture (I_{-1}) matters most (Fig. 6e). This is compatible with long intervals engendering a promise to oneself to leave straight after the next capture. The interval preceding this (I_{-2}) also has a significant effect.

The significant effect of time of arrival in the patch (t_a) implies that the policy was still adjusting even after the first 15 min. This was in the adaptive direction of shortening residence times: switching rate is 1.6 times higher in a pond entered after 45 min than that after 15 min.

Table 5. Cox regression model describing how various factors affect tendency to switch ponds

Variable	Transformation	Coefficient	P value	
Factors operating ≥ 4 s after a capture (otherwise values set to 0)				
I =interval since last capture	\sqrt{I}	0.266	<0.0001	} $R_0=2.38$ replaces when $N=0$
I_{-1} =interval preceding last capture	$\sqrt{I_{-1}}$	0.289	<0.0001	
$R=I/I_{-1}$: $R^+=1$ if $R \leq 1$ else $R^+=R$	$\ln[R^+]$	0.615	<0.0001	
Factors operating < 4 s after a capture (otherwise values set to 0)				
I_{-1} =interval preceding last capture	$\sqrt{I_{-1}}$	0.500–0.633	<0.0001	} $J_1=4.04$ replaces when $N=1$
I_{-2} =interval preceding I_{-1}	$\sqrt{I_{-2}}$	0.169	<0.0001	
Additional factors				
t_a =arrival time at pond	t_a	2.60×10^{-4}	<0.0001	
Subject (random factor)			<0.0001	

Positive coefficients indicate an increased tendency to switch: the exponent of the sum of the terms in the table is assumed to have a multiplicative effect on the baseline probability of switching, which is a fitted function of time at the pond (Fig. 6a). A range of coefficients is given for I_{-1} when $I < 4$ s because they differ ($P < 0.0001$) across environments. Units are seconds. The analysis omits ponds visited in the first 15 min.

Adaptive policy adjustment was further tested by including a factor describing rate of capture at the previous pond (captures/time spent). We expect high success rates to suggest a high return rate in the environment in general and thus to increase the tendency to leave the next patch. The relationship was just significant ($P = 0.043$) but would be more so without cases when the previous patch was left without a capture ($N_S/T_S = 0$ in Fig. 6f). A subject willing to leave one patch without a capture is liable to retain a similar strategy and leave the next patch relatively early also (more liable than at another time). So the proposed effect of success in one patch on the strategy used in the next is confounded by autocorrelation in the strategies applied.

We tested for a difference between environments by entering a four-level environment factor. Overall tendency to leave did not differ significantly between environments (mirroring the nonsignificant differences in T_S reported above) nor did the interactions with the main variables listed in Table 5, with the exception of the effect of I_{-1} on switching within 4 s of a capture ($P < 0.0001$). In the two aggregated environments, long periods without a capture increased the chances of switching directly after the next capture less than that in the more dispersed environments; for an intercapture interval of 45 s, the difference between environments in the leaving tendency immediately after the capture could be a factor of 2.5.

Importantly, there was no significant effect of the number of captures (N), whether we consider the situations within and after 4 s of a capture separately or combined ($P = 0.3, 0.9, 0.6$, respectively). When we include both N and an interaction of N with environment, the significance becomes barely marginal ($P = 0.1$); the coefficients suggest that any effect of N might exist only in the Poisson environment, where the least effect was predicted.

We checked for the validity of the proportional hazards model by examining whether the Schoenfeld residuals showed a relationship with T (Therneau 1999). Unfortunately there were significant effects for I and I_{-1} , both declining in influence at longer residence times rather than having a consistent effect. However, the decline appeared noticeable only with $T > 400$ s, when 95% of ponds have already been left; and using other transformations of I and

I_{-1} can remove the nonproportionality (at the expense of a poorer fit). We therefore consider the model in Table 5 a fair indication of how cues were used.

Self-reports

Eighty of the 100 subjects wrote about their strategies. One of us and one assistant independently scored whether each description mentioned any of three cues: fish caught from the current pond (N), time at the current pond (T) or interval without a capture (I). Descriptions were scored in random order, with scorers blind to the environment faced. The scorers disagreed whether the subjects had mentioned use of N in four of the 80 subjects, use of T in 16 subjects and use of I in 14 subjects. Nevertheless the pattern is clear: for 20 of 80 subjects, at least one of the two scorers thought that the subject mentioned N as a cue, the value for T is 27, but for I it is a much higher 59 (a proportion of 0.74). For each cue, we tested for a difference between environments in whether it was mentioned, pooling the two aggregated environments. There was only a marginally significant difference in mentions of T ($P = 0.04$ without correcting for multiple comparisons): T was more often mentioned in the aggregated environments.

Following this open-ended question, subjects explicitly rated how often they used each of the same three cues (Table 6). Each cue was rated by majorities of the subjects

Table 6. Subject ratings of how often they used a cue in deciding when to switch

Cue	Response			
	Not at all	Rarely	Often	Very often
T	4	29	47	20
N	5	27	39	29
I	1	8	33	58

For each cue the numbers of subjects (of 100) choosing each response are given.

as used 'often' or 'very often', but I was much the most often reported used 'very often'. Of the 100 subjects, 49 scored I as more often used than N compared with only 16 scoring N as more often used than I . Also, 53 subjects scored I above T compared with 10 scoring T above I . Differences between environments were not significant.

DISCUSSION

What Patch-leaving Rules do Humans Use?

We fitted two models: the T_S vs N_S linear regression fitted a model based on one family of simple deterministic rules (but allowing for imprecision in time estimation); the Cox proportional hazards regression fitted a model based on a probabilistic tendency to switch affected by time at the pond and other factors. Although the multiplicative weighting of probabilities is perhaps unrealistic as a model of cognition, Cox regression has an advantage in allowing statistical comparison of different cues. This established that recent intervals without captures are much better predictors of switching than number of captures, so the T_S vs N_S regression is also probably invalid as a direct description of the cognitive mechanism. But this linear regression remains a useful description of behaviour and allows direct comparison with the optimal policy.

The between-subject variation in the slope of the switching line swamps the between-environment variation (Fig. 5), although the former might be reduced by observing subjects for longer times. Nevertheless it is striking how consistently subjects generated a positive T_S vs N_S slope (92 of 100). What can we conclude about subjects' decision rules? Despite the simplicity of the fixed- N rule and that it was effectively optimal in the even environment, clearly nobody used it. It is harder to rule out individuals using a fixed- T rule because there will be some error in estimating time, and residence times that by chance are longer will tend to be associated with more captures. However, we showed by simulation that this explanation is insufficient to explain the mean slopes in three of the environments. Additional evidence comes from considering the within-subject variation in residence time over the last two thirds of the experimental session. In all but seven subjects the coefficient of variation is above 0.2 (minimum = 0.10, median = 0.37), whereas, even without consciously counting, humans can replicate a duration with coefficients of about 0.15 (reviewed by Wearden 1991).

The positive correlation between residence time (T_S) and number of captures (N_S) could potentially be generated by either an N -dependent T threshold or a giving-up-time (i.e. I -dependent) rule (Fig. 1). Both sets of cues can estimate the recent rate of capture, so both may lead to similar decisions. There are many other ways to combine N , T , I and the preceding intercapture intervals (Roche et al. 1998); for instance the giving-up time could be modified by N and/or T (e.g. Kamil et al. 1988; Wildhaber et al. 1994; Nelson & Roitberg 1995) or the incremental effect of a capture on tendency to stay could

depend on I_{-1} (cf. Burger et al. 2006). Several studies that used Cox regression (like ours with h_0 a function of T) concluded that there was an effect of N on leaving tendency without testing whether the concomitant increase in average I could be responsible instead (e.g. Wajnberg et al. 2000, 2003; van Baaren et al. 2005). Our Cox regressions implied that I and T but not N had influences.

The self-reports also implied that I was more often a component of subjects' strategies than N . This is despite N appearing on screen, whereas subjects had to rely on their own perception of time. However, self-reports of decision-making processes are often an unreliable guide (Nisbett & Wilson 1977), including those in progressive-ratio experiments (Wanchisen et al. 1992; Hackenberg & Axtell 1993).

The Cox regression suggested that subjects combine the information from the current interval without a capture with the previous such interval. The average of two intervals potentially provides a less noisy estimate of current capture rate than reliance on just one interval. However, simulations of such a rule (not presented here) showed an advantage over a simple giving-up-time rule only in the aggregated environment, in which the simple version anyway performed close to maximal. There was no indication that still earlier intercapture intervals affected decisions. Haccou et al. (1991), studying parasitoids, found shorter giving-up times when recent intervals were longer, suggesting similar averaging. Recent intercapture intervals also seemed important in determining when pigeons, *Columba livia*, switched (Roche et al. 1996). Reliance on mainly the most recent interval to predict the next interval is a widespread phenomenon in contexts beyond patch switching (Shettleworth & Plowright 1992).

The ratio of the current interval without a capture to the preceding interval was also a highly significant predictor in the Cox regression. Let us suppose this to be a real phenomenon, although it could emerge simply because of nonadditivity in the parallel, positive effects of I and I_{-1} . If the current interval was less than the last interval, the ratio had no effect, perhaps because such a pattern is expected and thus uninformative. A similar proportionality to the preceding interval describes when pigeons start to peck in expectation of a reward (Higa et al. 1991) and indeed the ratio of I to I_{-1} , or to an average of earlier intercapture intervals, has been implicated in their switching rules (Roche et al. 1998). Similarly, starlings, *Sturnus vulgaris*, stop pecking for a reward after a constant multiple of the wait that they had been trained to expect (Kacelnik & Brunner 2002). The ratio cue seems a good way to detect a sudden deterioration in the environment, for instance if all prey hide. However, if the ratio cue were used alone, a luckily brief intercapture interval would be likely to trigger premature leaving before another capture occurred.

In summary, the Cox regression found that several cues affect the tendency to leave and that there is some interaction between them. This apparently does not support the hypothesis of a simple deterministic rule of thumb (Hutchinson & Gigerenzer 2005), which had been suggested by the simplicity of the optimal strategy if the environment is known and by the robustness of certain simple rules under uncertainty about the environment.

But it remains possible that different individuals use different simple rules and that the complexity suggested by the Cox regression comes from amalgamating data from different individuals.

Leaving directly after a capture

The Cox regression indicates that a capture increases the tendency to switch directly afterwards, which explains the bimodality of giving-up times (Fig. 3). In isolation this would lead to shorter T_S when N_S is large, and thus a switching line with negative slope, but the effect is overwhelmed by the influences of I and I_{-1} . We proposed that a long fruitless wait prompted subjects to promise themselves to leave after the next capture. The motivation to wait until then rather than switch immediately would come from an illogical reluctance to 'waste' time already 'invested' in waiting. This would be a form of the Concorde fallacy (Dawkins 1976, page 162). Arkes & Ayton (1999) argue that this fallacy is widespread in humans but rare or absent in other animals because humans formulate abstract rules like 'don't waste' and then overgeneralize them to inappropriate situations. Indeed we know of no similar finding from patch leaving in other species: other cases of leaving straight after a capture are explicable as an adaptive response to a deterministic reward schedule (e.g. Ydenberg 1984) or to a dispersed distribution of items over patches (e.g. Redhead & Tyler 1988). In the latter example, the rats' tendency to switch directly after a capture increased over the course of the experiment, which suggests exactly the sort of facultative adjustment of switching rule (in response to consistent patch qualities) that we were seeking.

There are further alternative explanations for switching directly after a capture. (1) Perhaps a capture prompts subjects to consider whether to switch, whereas otherwise they were mostly concentrating on spotting fish or day-dreaming. (2) Subjects may imagine themselves to be observed; a real angler commented that leaving after a fruitless interval looks to others like a misjudgement and thus may involve a loss of face as well as a feeling of dissatisfaction. (3) In many foraging situations the timing of captures within a patch exhibits some regularity of spacing, in which case leaving after a capture can be adaptive. This regularity might arise because items are physically spaced apart or because catching one item makes other prey hide for some period.

Lack of Adjustment of Strategy between Environments

Our main prediction was that different dispersions of fish among ponds would prompt strategies differing dramatically in how captures affected switching. In the even environment a capture should encourage switching, in the Poisson environment it should have no effect and in the aggregated environments it should discourage switching. In fact for nearly all subjects N_S and T_S were positively correlated, with no significant effect of environment. The Cox regression found a significant effect of environment only on the extent to which longer intervals

triggered switching directly after a capture. Nor did the self-reports of cue use find convincing differences between environments.

This insensitivity might reflect the difficulty of learning the dispersion or the appropriate strategy while trying to exploit patches efficiently. A model of learning environmental quality has confirmed that stable estimates of variance in patch quality are obtainable rather later than accurate estimates of the mean (Rodríguez-Gironés & Vásquez 1997, although their learning procedure is suboptimal in several respects). Our intuition was that subjects might not be conscious of being in an even environment but that they should detect the unevenness of the aggregated environment once they encounter a very poor patch. A later study using a modified version of the experiment, in which subjects experienced an even and an aggregated environment in sequence, has supported this intuition by asking subjects explicitly about their perceptions of the dispersions (A. Wilke & R. Mata, unpublished data). In the aggregated environment with a 25-s travel time, our intuition is compatible with the change of T_S vs N_S slope during the course of the experimental session: the slope started off appropriate to a less-aggregated environment (but more aggregated than Poisson), and over the first 15 min it changed to one appropriate to the actual environment encountered (Fig. 4a). However, subjects experiencing the Poisson environment exhibited the same direction of change, which here is the opposite of that predicted.

One adaptive explanation for captures increasing the tendency to stay whatever the environment is that this policy is the best in an aggregated environment, that aggregated is the predominant pattern of dispersion in nature (e.g. Taylor et al. 1978) and thus that evolution has hardwired us to gamble on this pattern given the difficulties of estimating dispersion (Rodríguez-Gironés & Vásquez 1997). As a supporting argument, the costs of misjudging an even environment as aggregated are lower than vice versa (Table 2). Using I as a cue instead of N limits the inefficiency of the policy should it be the wrong bet.

However, Rodríguez-Gironés & Vásquez (1997) point out that many of the data claimed as evidence for aggregated distributions comprise numbers per randomly placed quadrat rather than per patch and thus may reflect merely patchiness itself. Competition between prey, and often patch depletion by other predators, reduce variation in patch quality. Also, often in real life cues other than the timing of captures indicate patch quality reliably; what matters is the posterior distribution of patch qualities given the information from these cues, so that a decremental rule can then become optimal even in an aggregated environment (Shaltiel & Ayal 1998; Driessen & Bernstein 1999).

Position of the Switching Line and Performance

Average residence time was considerably too long. Our prior intuition was the opposite, that subjects would

switch early to take a break from concentrating on the screen. An overwhelming majority of animal studies have also found a tendency to switch patch later than optimal (Nonacs 2001). One explanation is that optimality models do not take into account extra risks and uncertainties associated with moving between patches (natural selection may have hardwired the animal to behave as if a travel risk always existed even if none exists in the laboratory). A similar explanation is that subjects expect the experiment to finish soon and wish to avoid wasting the last period waiting for the next patch (cf. Wajenberg et al. 2006). We modified our optimization program so that there was a constant risk of the experiment ending and the currency maximized was total fish collected, rather than rate of capture; for reasonable risks of interruption the effect on average T_S was slight (cf. Newman 1991; Nonacs 2001), and fears of interruption would have to be unbelievably overblown for predicted T_S to match observations (perceived half-lives between 14 and 36 s: cf. Stephens 2002).

Another reason to stay longer than predicted might be to learn more about the environment (e.g. the form of the gain curves), which our optimality model unrealistically assumes are known (Oaten 1977). We expect this effect to be strongest early in the experimental session, although some effect may persist if subjects suspect that the environment may change. The Cox regression found that the tendency to leave indeed increased even within the later part of the experimental session (similarly, average T_S decreased over this period, although not significantly: $P = 0.060$). But another explanation for this is subjects learning that shorter residence times increased performance. In support of this explanation, leaving tendency increased more following a higher success rate on the previous pond. This makes sense: you should switch earlier if your estimation of the richness of other patches has increased (McNamara & Houston 1985). Other animals indeed adjust residence time adaptively according to experience on the previous patch and most recent travel times (e.g. Roitberg et al. 1992; Visser et al. 1992; Cuthill et al. 1994; Tenteliet al. 2006).

A Different Type of Explanation: Minimizing Delay to Reinforcement

One proposed explanation of nonoptimal behaviour in progressive-ratio experiments is that the animals are

concerned with a different currency: they short-sightedly seek to minimize delays to the next few rewards rather than maximize long-term reward rate (e.g. Mazur & Vaughan 1987; Shull & Spear 1987). Cassini et al. (1990) went further, proposing minimization of the delay to the very next reward alone as a currency compatible with behaviour of guinea pigs, *Cavia porcellus*, in a patch-leaving experiment (although a follow-up study reverted to a rate-maximization model: Cassini et al. 1993). The disproportionate importance of delay to the very next reward has long been established in other choice contexts and has been associated with impulsivity or lack of self control (Stephens & McLinn 2003). The delay-minimization criterion predicts longer residence times, because much of the advantage of switching under a rate-maximization criterion comes from the several short intercapture intervals after the first capture and not just from the shorter interval until that first capture (cf. the proof in Houston 1987 for deterministic schedules).

We used dynamic programming to calculate the strategy minimizing delay to next reward in each of our four environments. For the switch option the expected delay is the travel time plus the expected time to a capture when arriving at a new pond of unknown quality. The average residence times (T_S) predicted are now much closer to those observed (Table 7). Other signs of this good match are that, if we simulate these rules assuming the timing of captures experienced in each pond of the experiment, roughly half the ponds are left earlier and half later than actually observed (0.53:0.47), and the proportion of subjects that left mostly earlier than predicted does not differ significantly from the proportion leaving mostly later (Table 7).

Minimizing delay to reward still predicts T_S vs N_S slopes similar to those maximizing reward rate (cf. Tables 4 and 7). This is because both expected delay to reward and expected rate of reward covary closely with the Bayesian estimate of the number of fish remaining given N and T . Thus predicted slopes in the even and Poisson environments remain in conflict with the observations. Our earlier explanations, such as positive slopes being appropriate in the commonest type of environment, still potentially apply. But alternatively we propose that humans might apply general-purpose learning rules based on recent delays to reinforcement to compare the options of switching or staying; subjects would estimate the next delay to reward should they stay from the most recent

Table 7. Testing predictions supposing that expected time to the next capture is minimized

Environment	T_S vs N_S slope (s/fish)		Average T_S (s)		% Ponds left later than predicted	Subjects leaving most ponds earlier (vs later) than predicted
	Minimizing delay	Observed	Minimizing delay	Observed		
Even	-247	23.3	181	168	41	14 vs 8
Poisson	3.0	19.0	188	183	43	12 vs 13
Aggregated $v=25$ s	20.5	19.6	182	207	52	10 vs 13
Aggregated $v=15$ s	20.0	20.1	133	170	53	7 vs 18

Technical details as in Table 4. The proportions shown in the last column do not differ significantly between environments ($\chi^2_3 = 6.1, P = 0.4$) and the overall ratio (43:52) does not differ significantly from 1:1 (binomial test: $P = 0.1$). Analyses omit ponds visited in the first 15 min.

intercapture intervals on the current patch (rather than from N and T). The Cox regression implies that the interval since the last capture (I) is also involved as a ready and appropriate estimate of time until the next capture. In any case the important point is that extrapolating an estimate of the next delay to reward from recent delays in the current patch will always lead to a positive T_S vs N_S slope (e.g. giving-up-time rule in Table 4): the more captures in a set time, the shorter the intervals and consequently the longer the stays. So the new explanation for why we do not observe fixed- N , decremental or even fixed- T rules in environments where they would be optimal is that to compare the two options of staying or leaving humans apply a general-purpose learning rule that can produce only win-stay behaviour.

We find the evidence for short-sighted delay minimization as an explanation of our results no more than suggestive. Experiments involving a wider range of parameter values are needed to establish that the agreement of observation with prediction with respect to long residence times is not coincidental. Another issue is that delay minimization does not fit existing data from operant experiments on when to switch from a progressive-interval or progressive-ratio option: maximization of reward rate generated the best fit in some species, including humans, and, where not, what appeared to be minimized was the harmonic mean of the delays to the next four or so reinforcements (Houston 1987; Hackenberg & Axtell 1993); such a currency produces behaviour intermediate between delay minimization and reward maximization. It would be surprising that humans are more prone to minimize delay than other animals because the conditional reinforcers (e.g. money) used in this and most human experiments do not differentially reinforce obtaining rewards quickly, as might food. However, these progressive-ratio experiments on humans involved deterministic intervals: it would be fascinating if with stochastic intervals the extra difficulty of averaging success over multiple patches as well as within each patch causes humans to give up on that computation and consider only the interval to the next reward. Another difference is that in the progressive-ratio experiments it was behaviour after multiple training sessions that fitted rate maximizing; when first encountering a patch-leaving problem humans might dare to attempt only the cognitively simpler strategy of minimizing delay and later as they gain experience switch to maximizing rate.

Further Investigations and Applications of Human Patch Leaving

Informing human subjects about the mechanics of this experiment was much quicker than training other animals for a similar operant task. Consequently we achieved subject sample sizes an order of magnitude greater than those in most studies of vertebrate patch leaving. This was valuable given the considerable interindividual variation encountered. It is also feasible to train a smaller number of humans over a longer period so that they can estimate environmental parameters better and hone strategies further (e.g. Wanchisen et al. 1992; Hackenberg & Axtell

1993). The limitation with humans is that it is unfeasible to train them in a closed economy without the distractions of similar tasks in everyday life.

Psychologists currently debate how modular the human brain is (Barrett & Kurzban 2006): one might expect that judgement of when to switch from a deteriorating resource would be an ideal candidate for such a module, given its likely importance throughout our evolutionary history (cf. Cosmides & Tooby 1994; Hills 2006). We think our cover story and animation would be sufficiently realistic to invoke such a module. Furthermore, unlike in many other studies, our depletion curves were realistic: intervals were stochastic with the appearance rate proportional to the number of prey remaining. Nevertheless it might be worthwhile to replicate the experiment using a hands-on task in which foraging and travelling involve more effort than watching a monitor. For instance, Miller (1999) describes a task involving searching for beans in buckets of rice.

However, the skill assayed in our experiment is applicable far beyond physical foraging tasks, to all sorts of task switching in the modern world, and not only to seeking information in the external environment (Pirolli & Card 1999) but also to searching in memory and problem solving. For instance we have run a similar set of experiments in which the task is to construct words out of a sequence of letters; finding solutions grows harder, so at some point subjects should request a new letter sequence (Wilke et al. 2004; Wilke 2006). When to switch from a deteriorating resource can also be viewed as a special case of the exploitation–exploration trade-off (Cohen et al. 2007).

Subsequent research should test whether humans ever apply fixed- N or fixed- T strategies, for instance even if explicitly informed that initial patch quality does not vary. Another project would be to manipulate the timing of fish appearances so as to quantify more exactly how information about recent intervals is combined. Such experiments make sense only when applied to subjects practised enough that their strategies are stable; detailed knowledge of a few subjects seems also the best way to establish the extent and nature of interindividual differences in strategy.

Acknowledgments

We thank Gregor Caregnato for organizing subjects and running the experiment, Julian Marewski for testing the program and other technical help, Christian Elsner for excellent IT support and Innes Cuthill, Thomas Hills, Alasdair Houston, Rui Mata, Jeff Stevens and an anonymous referee for their thoughtful comments on the text. A particularly big thank you goes to Uwe Czienskowski for programming the experimental interface. A.W. carried out the empirical work and initial analysis as part of his doctoral dissertation (Wilke 2006). He has been subsequently supported by Research Fellowship WI 3215/1-1 from the Deutsche Forschungsgemeinschaft and the Owen F. Aldis Scholarship Fund from the International Society for Human Ethology. The Ethics Committee of the Max Planck Institute for Human Development approved the experiment.

Supplementary Material

Supplementary material for this article can be found with the online version at [doi:10.1016/j.anbehav.2007.09.006](https://doi.org/10.1016/j.anbehav.2007.09.006).

References

- Adler, F. R. & Kotar, M. 1999. Departure time versus departure rate: how to forage optimally when you are stupid. *Evolutionary Ecology Research*, **1**, 411–421.
- Agetsuma, N. 1999. Simulation of patch use by monkeys using operant conditioning. *Journal of Ethology*, **16**, 49–55.
- van Alphen, J. J. M., Bernstein, C. & Driessen, G. 2003. Information acquisition and time allocation in insect parasitoids. *Trends in Ecology & Evolution*, **18**, 81–87. doi:10.1016/S0169-5347(02)00035-6.
- Arkes, H. R. & Ayton, P. 1999. The sunk cost and concordance effects: are humans less rational than lower animals? *Psychological Bulletin*, **125**, 591–600.
- Aswani, S. 1998. Patterns of marine harvest effort in southwestern New Georgia, Solomon Islands: resource management or optimal foraging? *Ocean and Coastal Management*, **40**, 207–235. doi:10.1016/S0964-5691(98)00047-7.
- van Baaren, J., Outreman, Y. & Boivin, G. 2005. Effect of low temperature exposure on oviposition behaviour and patch exploitation strategy in parasitic wasps. *Animal Behaviour*, **70**, 153–163. doi:10.1016/j.anbehav.2004.10.016.
- Barrett, H. C. & Kurzban, R. 2006. Modularity in cognition: framing the debate. *Psychological Review*, **113**, 628–647. doi:10.1037/0033-295X.113.3.628.
- Brunner, D., Kacelnik, A. & Gibbon, J. 1992. Optimal foraging and timing processes in the starling, *Sturnus vulgaris*: effect of inter-capture interval. *Animal Behaviour*, **44**, 597–613. doi:10.1016/j.anbehav.2004.10.016.
- Burger, J. M. S., Huang, Y., Hemerik, L., van Lenteren, J. C. & Vet, L. E. M. 2006. Flexible use of patch-leaving mechanisms in a parasitoid wasp. *Journal of Insect Behavior*, **19**, 155–170. doi:10.1007/s10905-006-9014-7.
- Burger, O., Hamilton, M. J. & Walker, R. 2005. The prey as patch model: optimal handling of resources with diminishing returns. *Journal of Archaeological Science*, **32**, 1147–1158. doi:10.1016/j.jas.2005.02.012.
- Cassini, M. H., Kacelnik, A. & Segura, E. T. 1990. The tale of the screaming hairy armadillo, the guinea pig and the marginal value theorem. *Animal Behaviour*, **39**, 1030–1050. doi:10.1016/S0003-3472(05)80776-6.
- Cassini, M. H., Lichtenstein, G., Ongay, J. P. & Kacelnik, A. 1993. Foraging behaviour in guinea pigs: further tests of the marginal value theorem. *Behavioural Processes*, **29**, 99–112. doi:10.1016/0376-6357(93)90030-U.
- Charnov, E. L. 1976. Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, **9**, 129–136. doi:10.1016/0040-5809(76)90040-X.
- Cohen, J. D., McClure, S. M. & Yu, A. J. 2007. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London, Series B*, **362**, 933–942. doi:10.1098/rstb.2007.2098.
- Cosmides, L. & Tooby, J. 1994. Origins of domain specificity: the evolution of functional organization. In: *Mapping the Mind: Domain Specificity in Cognition and Culture* (Ed. by L. A. Hirschfeld & S. A. Gelman), pp. 85–116. New York: Cambridge University Press.
- Cuthill, I. C., Haccou, P. & Kacelnik, A. 1994. Starlings (*Sturnus vulgaris*) exploiting patches: response to long-term changes in travel time. *Behavioral Ecology*, **5**, 81–90. doi:10.1093/beheco/5.1.81.
- Czienskowski, U. 2005. *The Fishing Task Experiment*. [computer software]. Berlin: Max Planck Institute for Human Development.
- Dall, S. R. X. 1995. Information and the generalist predator problem. Ph.D. thesis, University of Bristol.
- Dawkins, R. 1976. *The Selfish Gene*. Oxford: Oxford University Press.
- Driessen, G. & Bernstein, C. 1999. Patch departure mechanisms and optimal host exploitation in an insect parasitoid. *Journal of Animal Ecology*, **68**, 445–459. doi:10.1046/j.1365-2656.1999.00296.x.
- van Gils, J. A., Schenk, I. W., Bos, O. & Piersma, T. 2003. Incompletely informed shorebirds that face a digestive constraint maximize net energy gain when exploiting patches. *American Naturalist*, **161**, 777–793.
- Goubault, M., Outreman, Y., Poinso, D. & Cortesero, A. M. 2005. Patch exploitation strategies of parasitic wasps under intra-specific competition. *Behavioral Ecology*, **16**, 693–701. doi:10.1093/beheco/ari043.
- Green, R. F. 1980. Bayesian birds: a simple example of Oaten's stochastic model of optimal foraging. *Theoretical Population Biology*, **18**, 244–256. doi:10.1016/0040-5809(80)90051-9.
- Green, R. F. 1984. Stopping rules for optimal foragers. *American Naturalist*, **123**, 30–43.
- Green, R. F. 1987. Stochastic models of optimal foraging. In: *Foraging Behavior* (Ed. by A. C. Kamil, J. R. Krebs & H. R. Pulliam), pp. 273–302. New York: Plenum.
- Green, R. F. 2006. A simpler, more general method of finding the optimal foraging strategy for Bayesian birds. *Oikos*, **112**, 274–284. doi:10.1111/j.0030-1299.2006.13462.x.
- Haccou, P., de Vlas, S. J., van Alphen, J. J. M. & Visser, M. E. 1991. Information processing by foragers: effects of intra-patch experience on the leaving tendency of *Leptopilina heterotoma*. *Journal of Animal Ecology*, **60**, 93–106.
- Hackenberg, T. D. & Axtell, S. A. M. 1993. Humans' choices in situations of time-based diminishing returns. *Journal of the Experimental Analysis of Behavior*, **59**, 445–470. doi:10.1901/jeab.1993.59-445.
- Higa, J. J., Wynne, C. D. L. & Staddon, J. E. R. 1991. Dynamics of time discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*, **17**, 281–291.
- Hills, T. T. 2006. Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, **30**, 3–41. doi:10.1207/s15516709cog0000_50.
- Houston, A. I. 1987. The control of foraging decisions. In: *Quantitative Analyses of Behavior: Vol. 6. Foraging* (Ed. by M. L. Commons, A. Kacelnik & S. J. Shettleworth), pp. 41–61. Hillsdale, New Jersey: L. Erlbaum.
- Hutchinson, J. M. C. & Gigerenzer, G. 2005. Simple heuristics and rules of thumb: where psychologists and behavioural biologists might meet. *Behavioural Processes*, **69**, 97–124. doi:10.1016/j.beproc.2005.02.019.
- Iwasa, Y., Higashi, M. & Yamamura, N. 1981. Prey distribution as a factor determining the choice of optimal foraging strategy. *American Naturalist*, **117**, 710–723.
- Kacelnik, A. & Brunner, D. 2002. Timing and foraging: Gibbon's scalar expectancy theory and optimal patch exploitation. *Learning & Motivation*, **33**, 177–195. doi:10.1006/lmot.2001.1110.
- Kamil, A. C. & Clements, K. C. 1990. Learning, memory, and foraging behavior. In: *Contemporary Issues in Comparative Psychology* (Ed. by D. A. Dewsbury), pp. 7–30. Sunderland, Massachusetts: Sinauer.
- Kamil, A. C. & Yoerg, S. I. 1985. The effects of prey depletion on the patch choice of foraging blue jays (*Cyanocitta cristata*). *Animal Behaviour*, **33**, 1089–1095. doi:10.1016/S0003-3472(85)80167-6.

- Kamil, A. C., Yoerg, S. I. & Clements, K. C. 1988. Rules to leave by: patch departure in foraging blue jays. *Animal Behaviour*, **36**, 843–853. doi:10.1016/S0003-3472(88)80167-2.
- Klaasen, R. H. G., Nolet, B. A. & van Leeuwen, C. H. A. 2007. Prior knowledge about spatial pattern affects patch assessment rather than movement between patches in tactile-feeding mallard. *Journal of Animal Ecology*, **76**, 20–29. doi:10.1111/j.1365-2656.2006.01184.x.
- Kramer, D. L. & Weary, D. M. 1991. Exploration versus exploitation: a field study of time allocation to environmental tracking by foraging chipmunks. *Animal Behaviour*, **41**, 443–449. doi:10.1016/S0003-3472(05)80846-2.
- McNair, J. M. 1982. Optimal giving-up times and the marginal value theorem. *American Naturalist*, **119**, 511–529.
- McNamara, J. M. & Houston, A. I. 1985. Optimal foraging and learning. *Journal of Theoretical Biology*, **117**, 231–249. doi:10.1016/S0022-5193(85)80219-8.
- Marschall, E. A., Chesson, P. L. & Stein, R. A. 1989. Foraging in a patchy environment: prey-encounter rate and residence time distributions. *Animal Behaviour*, **37**, 444–454. doi:10.1016/0003-3472(89)90091-2.
- Mazur, J. E. & Vaughan, W., Jr. 1987. Molar optimization versus delayed reinforcement as explanations of choice between fixed-ratio and progressive-ratio schedules. *Journal of the Experimental Analysis of Behavior*, **48**, 251–261. doi:10.1901/jeab.1987.48-251.
- Metcalfe, D. & Barlow, K. R. 1992. A model for exploring the optimal trade-off between field processing and transport. *American Anthropologist*, **94**, 340–356.
- Miller, T. E. 1999. Optimal foraging: balancing costs and rewards in making foraging decisions. In: *Tested Studies for Laboratory Teaching*, Vol. 20 (Ed. by S. J. Karcher), pp. 187–205. Proceedings of the 20th Workshop/Conference of the Association for Biology Laboratory Education.
- Nelson, J. M. & Roitberg, B. D. 1995. Flexible patch time allocation by the leafminer parasitoid, *Opius dimidiatus*. *Ecological Entomology*, **20**, 245–252.
- Newman, J. A. 1991. Patch use under predation hazard: foraging behavior in a simple stochastic environment. *Oikos*, **61**, 29–44.
- Nisbett, R. E. & Wilson, T. D. 1977. Telling more than we can know: verbal reports on mental processes. *Psychological Review*, **84**, 231–259.
- Nonacs, P. 2001. State dependent behavior and the marginal value theorem. *Behavioral Ecology*, **12**, 71–83.
- Oaten, A. 1977. Optimal foraging in patches: a case for stochasticity. *Theoretical Population Biology*, **12**, 263–285. doi:10.1016/0040-5809(77)90046-6.
- Ollason, J. G. 1980. Learning to forage—optimally. *Theoretical Population Biology*, **18**, 44–56. doi:10.1016/0040-5809(80)90039-8.
- Olsson, O. & Brown, J. S. 2006. The foraging benefits of information and the penalty of ignorance. *Oikos*, **112**, 260–273. doi:10.1111/j.0030-1299.2006.13548.x.
- Olsson, O. & Holmgren, N. M. A. 1998. The survival-rate-maximizing policy for Bayesian foragers: wait for good news. *Behavioral Ecology*, **9**, 345–353. doi:10.1093/beheco/9.4.345.
- Outreman, Y., Le Ralec, A., Wajnberg, E. & Pierre, J.-S. 2005. Effects of within- and among-patch experiences of the patch-leaving decision rules in an insect parasitoid. *Behavioral Ecology and Sociobiology*, **58**, 208–217. doi:10.1007/s00265-004-0895-1.
- Pierre, J.-S., van Baaren, J. & Boivin, G. 2003. Patch leaving decision rules in parasitoids: do they use sequential decisional sampling? *Behavioral Ecology and Sociobiology*, **54**, 147–155. doi:10.1007/s00265-003-0617-0.
- Pinheiro, J. C. & Bates, D. M. 2000. *Mixed-effects Models in S and S-PLUS*. New York: Springer.
- Pirolli, P. 2005. Rational analyses of information foraging on the web. *Cognitive Science*, **29**, 343–373. doi:10.1207/s15516709cog0000_20.
- Pirolli, P. & Card, S. 1999. Information foraging. *Psychological Review*, **106**, 643–675.
- Redhead, E. & Tyler, P. A. 1988. An experimental analysis of optimal foraging behaviour in patchy environments. *Quarterly Journal of Experimental Psychology*, **40B**, 83–102.
- Roche, J. P. 1996. Patch-leaving decisions in black-capped chickadees. *Animal Behaviour*, **52**, 289–298. doi:10.1006/anbe.1996.0174.
- Roche, J. P., Stubbs, D. A. & Glanz, W. E. 1996. Assessment and choice: an operant simulation of foraging in patches. *Journal of the Experimental Analysis of Behavior*, **66**, 327–347. doi:10.1901/jeab.1996.66-327.
- Roche, J. P., Timberlake, W., Glanz, W. E. & Stubbs, D. A. 1998. The influence of current-visit experience within a prey patch on patch persistence. *Behavioural Processes*, **43**, 11–25. doi:10.1016/S0376-6357(97)00083-1.
- Rodríguez-Gironés, M. A. & Vásquez, R. A. 1997. Density-dependent patch exploitation and acquisition of environmental information. *Theoretical Population Biology*, **52**, 32–42. doi:10.1006/tubi.1997.1317.
- Roitberg, B. D., Mangel, M., Lalonde, R. G., Roitberg, C. A., van Alphen, J. J. M. & Vet, L. 1992. Seasonal dynamic shifts in patch exploitation by parasitic wasps. *Behavioral Ecology*, **3**, 156–165. doi:10.1093/beheco/3.2.156.
- Roitberg, B. D., Sircom, J., Roitberg, C. A., van Alphen, J. J. M. & Mangel, M. 1993. Life expectancy and reproduction. *Nature*, **364**, 108. doi:10.1038/364108a0.
- Shaltiel, L. & Ayal, Y. 1998. The use of kairomones for foraging decisions by an aphid parasitoid in small host aggregations. *Ecological Entomology*, **23**, 319–329. doi:10.1046/j.1365-2311.1998.00141.x.
- Shettleworth, S. J. & Plowright, C. M. S. 1992. How pigeons estimate rates of prey encounter. *Journal of Experimental Psychology: Animal Behavior Processes*, **18**, 219–235.
- Shull, R. L. & Spear, D. J. 1987. Detention time after reinforcement: effects due to delay of reinforcement? In: *Quantitative Analyses of Behavior*, Vol. 5. *The Effect of Delay and of Intervening Events on Reinforcement Value* (Ed. by M. L. Commons, J. E. Mazur, J. A. Nevin & H. Rachlin), pp. 187–204. Hillsdale, New Jersey: L. Erlbaum.
- Smith, E. A. 1983. Anthropological applications of optimal foraging theory: a critical review. *Current Anthropology*, **24**, 625–651.
- Smith, E. A. 1991. *Inujjamiut Foraging Strategies: Evolutionary Ecology of an Arctic Hunting Economy*. New York: Aldine de Gruyter.
- Smith, E. A. & Wishnie, M. 2000. Conservation and subsistence in small-scale societies. *Annual Review of Anthropology*, **29**, 493–524. doi:10.1146/annurev.anthro.29.1.493.
- Stephens, D. W. 2002. Discrimination, discounting and impulsivity: a role for an informational constraint. *Philosophical Transactions of the Royal Society of London, Series B*, **357**, 1527–1537. doi:10.1098/rstb.2002.1062.
- Stephens, D. W. & Krebs, J. R. 1986. *Foraging Theory*. Princeton: Princeton University Press.
- Stephens, D. W. & McLinn, C. M. 2003. Choice and context: testing a simple short-term choice rule. *Animal Behaviour*, **66**, 59–70. doi:10.1006/anbe.2003.2177.
- Taylor, L. R., Woiwod, I. P. & Perry, J. N. 1978. The density-dependence of spatial behaviour and the rarity of randomness. *Journal of Animal Ecology*, **47**, 383–406.
- Tenhumberg, B., Keller, M. A., Possingham, H. P. & Tyre, A. J. 2001. Optimal patch-leaving behaviour: a case study using the

- parasitoid *Cotesia rubecula*. *Journal of Animal Ecology*, **70**, 683–691. doi:10.1046/j.1365-2656.2001.00530.x.
- Tentelier, C., Desouhant, E. & Fauvergue, X.** 2006. Habitat assessment by parasitoids: mechanisms for patch use behavior. *Behavioral Ecology*, **17**, 515–521. doi:10.1093/beheco/arj062.
- Therneau, T. M.** 1999. A package for survival analysis in S. Technical Report No. 53, Mayo Clinic, Rochester, MN <http://mayoresearch.mayo.edu/mayo/research/biostat/upload/survival.pdf>.
- Therneau, T. M., Grambsch, P. M. & Pankratz, V. S.** 2003. Penalized survival models and frailty. *Journal of Computational and Graphical Statistics*, **12**, 156–175. doi:10.1198/1061860031365.
- Visser, M. E., van Alphen, J. J. M. & Nell, H. W.** 1992. Adaptive superparasitism and patch time allocation in solitary parasitoids: the influence of pre-patch experience. *Behavioral Ecology and Sociobiology*, **31**, 163–171. doi:10.1007/BF00168643.
- Vos, M., Hemerik, L. & Vet, L. E. M.** 1998. Patch exploitation by the parasitoids *Cotesia rubecula* and *Cotesia glomerata* in multi-patch environments with different host distributions. *Journal of Animal Ecology*, **67**, 774–783. doi:10.1046/j.1365-2656.1998.00239.x.
- Waage, J. K.** 1979. Foraging for patchily-distributed hosts by the parasitoid, *Nemeritis canescens*. *Journal of Animal Ecology*, **48**, 353–371.
- Wajnberg, E.** 2006. Time allocation strategies in insect parasitoids: from ultimate predictions to proximate behavioral mechanisms. *Behavioral Ecology and Sociobiology*, **60**, 589–611. doi:10.1007/s00265-006-0198-9.
- Wajnberg, E., Rosi, M. C. & Colazza, S.** 1999. Genetic variation in patch time allocation in a parasitic wasp. *Journal of Animal Ecology*, **68**, 121–133. doi:10.1046/j.1365-2656.1999.00270.x.
- Wajnberg, E., Fauvergue, X. & Pons, O.** 2000. Patch leaving decision rules and the marginal value theorem: an experimental analysis and a simulation model. *Behavioral Ecology*, **11**, 577–586. doi:10.1093/beheco/11.6.577.
- Wajnberg, E., Gonsard, P.-A., Tabone, E., Curty, C., Lezcano, N. & Colazza, S.** 2003. A comparative analysis of patch-leaving decision rules in a parasitoid family. *Journal of Animal Ecology*, **72**, 618–626. doi:10.1046/j.1365-2656.2003.00731.x.
- Wajnberg, E., Bernhard, P., Hamelin, F. & Boivin, G.** 2006. Optimal patch time allocation for time-limited foragers. *Behavioral Ecology and Sociobiology*, **60**, 1–10. doi:10.1007/s00265-005-0131-7.
- Wanchisen, B. A., Tatham, T. A. & Hineline, P. N.** 1992. Human choice in “counterintuitive” situations: fixed- versus progressive-ratio schedules. *Journal of the Experimental Analysis of Behavior*, **58**, 67–85. doi:10.1901/jeab.1992.58-67.
- Wearden, J. H.** 1991. Do humans possess an internal clock with scalar timing properties? *Learning & Motivation*, **22**, 59–83.
- Wildhaber, M. L., Green, R. F. & Crowder, L. B.** 1994. Bluegills continuously update patch giving-up times based on foraging experience. *Animal Behaviour*, **47**, 501–513. doi:10.1006/anbe.1994.1075.
- Wilke, A.** 2006. Evolved responses to an uncertain world. Ph.D. thesis, Free University of Berlin.
- Wilke, A., Hutchinson, J. M. C. & Todd, P. M.** 2004. Testing simple rules for human foraging in patchy environments. In: *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society* (Ed. by K. Forbus, D. Gentner & T. Regier), pp. 1656. Mahwah, New Jersey: L. Erlbaum.
- Ydenberg, R. C.** 1984. Great tits and giving-up times: decision rules for leaving patches. *Behaviour*, **90**, 1–24.