Paper 3

Sigrunn Eliassen, Christian Jørgensen & Jarl Giske. Exploration or exploitation: life expectancy changes the value of learning in foraging strategies. – *Manuscript conditionally accepted in Oikos*

Exploration or exploitation: life expectancy changes the value of learning in foraging strategies

Sigrunn Eliassen, Christian Jørgensen and Jarl Giske

Department of Biology, University of Bergen, P. O. Box 7800, N-5020 Bergen, Norway

Abstract

The acquisition of information is a fundamental part of individual foraging behaviour in heterogeneous and changing environments. We examine how foragers may benefit from utilizing a simple learning rule to update estimates of temporal changes in resource levels. In the model, initial expectation of resource conditions and rate of replacing past information by new experiences are genetically inherited traits. Patch-time allocation differs between learners and foragers that use a fixed patch-leaving threshold throughout the foraging season. It also deviates from foragers that obtain information about the environment at no cost. At the start of a foraging season, learners sample the environment by frequent movements between patches, sacrificing current resource intake for information acquisition. This is done to obtain more precise and accurate estimates of resource levels, resulting in increased intake rates later in season. Risk of mortality may alter the trade-off between exploration and exploitation and thus change patch sampling effort. As lifetime expectancy decreases, learners invest less in information acquisition and show lower foraging performance when resource level changes through time.

Introduction

Foragers that explore various prey types or resource patches may obtain information that can enhance future foraging performance. In such learning processes, foragers integrate pieces of information acquired through experience (Stephens, 1993). Learners may adjust their behaviour to changing environments when fitness consequences of a given action vary within an individual's lifetime.

Information comes at a cost, however, as sampling often conflicts with resource exploitation or other activities (e.g. Stephens, 1987; Krebs & Inman, 1992). Traditional models of patch use (Charnov, 1976) and forager distributions (Fretwell & Lucas, 1970) commonly assume that animals have complete and free information about the spatial and temporal distribution of resources. When information has to be actively sampled and used, foraging behaviour often deviates from patterns predicted for such omniscient foragers (Abrahams, 1986; Bernstein *et al.*, 1988; Rodriguez-Gironés & Vásquez, 1997).

When resource conditions vary temporally, new experiences become more valuable than older information, and a foragers should bias its estimate towards recent information (McNamara & Houston, 1985, 1987). However, with more weight given to each sample, estimates becomes more sensitive to natural variability and sampling errors (McNamara & Houston, 1985; Hirvonen *et al.*, 1999). In changing environments, this imposes a trade-off between having a precise estimate on the one hand, and keeping the world-view up to date on the other.

Foraging models commonly consider only the ability of learners to reduce uncertainty in estimates of the environment (Mangel & Clark, 1983; Dall *et al.*, 2005). Information is, however, only valuable when knowledge can lead to changes in behaviour that have fitness consequences (Gould, 1974; Stephens, 1987; Dall *et al.*, 2005). This perspective has important implications when studying foraging behaviour in heterogeneous environments. If foragers cannot change actions, or if behavioural shifts

have small fitness consequences, then the value of information is low (Mangel, 1990). In such cases, there is no need to accurately estimate environmental parameters or to learn from experience.

Information is thus a two-sided coin. On the one side, it does not exist separately from the environment, but must be collected, often by altering behaviour. On the other side, information is only valuable as long as it leads to more efficient foraging behaviour. The benefit of information must therefore be traded against investment in other activities, for instance when there is a conflict between information and resource harvesting (Clark & Mangel, 1984). This does not necessarily mean that foraging and sampling are mutually exclusive behaviours. Commonly there is an interaction between the two (e.g. Cohen, 1993), for example when information on patch quality is gained from resource encounters and search times (Oaten, 1977; Iwasa *et al.*, 1981; McNamara, 1982; Green, 1984; Olsson & Holmgren, 1998). Foragers that explore several patches would gain a better estimate of the general resource level in the habitat, but frequent patch shifts could interfere with resource exploitation and the assessment made in each patch (Stephens, 1987; Valone, 1989; Mangel, 1990). The investment to improve future foraging performance thus manifests itself as reduced intake rates during intensive sampling periods.

Behaviours commonly range from those that are unaffected by specific experiences to those that wholly depend on them (Arak & Enquist, 1993; Papaj, 1993; Dukas, 1998). Evolutionarily adaptive foraging strategies should balance present and future foraging benefits, and make the best of both spatial and temporal heterogeneities. The precise nature of such adaptive strategies is far from obvious (Clark & Mangel, 1986). Hence, to predict the way optimal learning strategies change with environmental characteristics, we need to understand the trade-offs in behaviour and life-history that foragers face. We study adaptive foraging behaviour that is partly specified by genetically inherited traits and partly updated with knowledge acquired through the forager's lifetime. The acquisition of information is treated as an integral part of the patch-leaving behaviour, which emphasizes the trade-off between investment in

exploration (information harvesting) and exploitation (immediate foraging reward). We study how learning strategies reflect the trade-off between the rate of information updating and the resulting uncertainty in estimates. Costs and benefits of learning are dynamic properties that depend on local patch assessment, the updating mechanism, and characteristics of the resource environment. We aim at evaluating how risk of mortality may alter the trade-off between exploration and exploitation, and thus change patch sampling effort and the quality of resource estimates.

Model description

We study foragers that explore a patchy and temporally changing habitat. Individuals search for discrete food items that are distributed in well-defined patches. By consuming resources, foragers deplete the patch, and intake rate drops. Foragers record the time between resource encounters and leave the patch when the search time exceeds a giving-up threshold. *Learning Foragers* (also termed *Learners*) use their past experiences to estimate global resource conditions. The learning rule is simple, with a genetically determined learning factor that weights past information against present. We use a genetic algorithm to evolve learning factors and initial giving-up thresholds. To track temporal changes in average resource conditions, a learner may sample different resource patches in the habitat and update patch-leaving estimates from one patch to the next.

The costs and benefits of learning emerge from interactions between the foraging strategy and the environment. We compare behaviour and performance of *Learning Foragers* to:

1) *Informed Foragers* that have free access to information about the global resource level. They know the optimal giving-up threshold and how it responds to seasonal changes in resource conditions.

2) *Fixed Foragers* that use a single genetically determined giving-up threshold throughout the foraging season, thereby ignoring information about temporal changes in resource levels.

All three strategies have the same local assessment problem, and differ only in the way they access and utilize information about global resource conditions. We compare *Learning Foragers* to *Fixed Foragers* to determine the value of using environmental information and to *Informed Foragers* in order to evaluate the costs of uncertainty. We analyze the emergent costs and benefits of different learning rules to see how sampling is balanced against resource exploitation in different environments.

Resource encounters

The lifetime of a forager is divided into discrete time steps, in which the forager either searches for resources within a patch or moves between resource patches. Resource density in the habitat changes temporally with a constant factor g every time step:

$$R_t = R_0 + gt$$

R_0 is the initial resource density and t denotes the number of time steps since the start of a foraging season. We assume that a forager never returns to previously exploited patches and that there is no competition for resources within a patch. The number of food items remaining in the current patch at time t is therefore given by:

$$r_t = R_t - k$$

where k is the number of resources a forager has encountered in that patch. The forager searches randomly for food within a patch and each time step, the probability P_e that it will encounter a resource item depends on r_t and the search efficiency a:

$$P_e = 1 - e^{-ar_t}$$

Whenever a uniformly distributed random number exceeds P_e the forager will hence encounter one resource. A forager can handle no more than one resource each time step.

Patch assessment and the patch-departure rule

When a forager consumes resources, the patch is gradually depleted and the encounter rate drops. A forager uses time since last resource encounter to estimate the current quality of a patch. The encounter process is, however, stochastic in nature, and consequently this estimate is associated with uncertainty (Iwasa *et al.*, 1981; Green, 1984). The patch-departure rule is inspired by the Marginal Value Theorem (Charnov, 1976) where a forager leaves a patch when resource intake rate falls below the environmental average. Hence, when time spent searching for the next resource exceeds a giving-up threshold, the forager will leave the patch to look for a new feeding location. To locate a new, randomly selected patch, the forager needs to travel for a fixed number of time steps *d*.

Learning

A *Learning Forager* may update its estimate of the environment through experiences made in previously visited patches. Upon leaving a patch, a forager calculates the average time between resource encounters, including travel time *d*:

$$\bar{n}_p = \frac{d + \sum_{i=1}^{k} n_i}{k}$$

Here, n_i is the time searched before finding the *i*'th of the total *k* resources encountered in patch *p*. Learning is incorporated into the model by the use of a temporal weighting rule, where the new giving-up threshold τ_{p+1} is a weighted average of the former threshold τ_p and experiences made in the last patch:

$$\tau_{p+1} = (1 - \gamma)\tau_p + \gamma\bar{n}_p$$

The learning factor γ determines the relative weighting of past information and the new patch quality sample. The learning rule is a linear operator (McNamara &

Houston, 1987; Mangel, 1990) and like Bayesian updating rules it integrates prior estimates and new data.

Evolving behavioural strategies

The initial giving-up threshold τ_0 and the learning factor γ can be regarded as genetically inherited traits. The learning factor does not change during the lifetime of a forager, but the initial giving-up threshold τ_0 is updated by experience. We find the best combination of learning factors and initial giving-up thresholds using a genetic algorithm (Sumida *et al.*, 1990; Holland, 1992). The learning factor may take any value between 0.0 and 1.0, and initial giving-up thresholds range from 1.0-50.0. Each model run was initiated with a population of 10 000 foragers, having randomly assigned values of τ_0 and γ. We evaluated the performance of the patch-leaving strategies by simulating individual foraging behaviour under different environmental conditions. Mortality is a stochastic process in the model, and all individuals have an equal probability P_s of surviving each time-step until the end of season T when all foraging activity ceases. Mortality rate does not change between years and we only consider within-year variation in resource levels.

In order to produce an offspring, a forager needs to accumulate a certain amount of resources. A constant fraction of the offspring survives and produces a new population of foragers that enter the foraging habitat next season. The fitness of a forager is hence proportional to its offspring production, which in turn depends on the amount of resources accumulated during a lifetime.

The genetic strings are haploid in the model and trait values are coded as continuous numbers. An offspring inherits both the learning factor and the initial giving-up threshold from its parent, unless the genetic traits are recombined with a randomly chosen partner. Recombination occurs with a probability M_r and then the offspring inherits one trait from each parent. There is a probability M_m that mutations may alter a trait and change its initial value by some random fraction M_s. We evolved the learning

rule traits for 3000 generations and used population means of τ_0 and γ when comparing foraging performance between strategies and resource environments.

Informed Foragers and the cost of uncertainty

An *Informed Forager* knows the best giving-up threshold to select when it enters a new patch. We used dynamic programming (Houston & McNamara, 1999; Clark & Mangel, 2000) to calculate the optimal $\tau^*(t)$ that maximizes the expected future reproductive output for a forager entering a patch at time t. Computational details are presented in Appendix 1 along with dynamic programming equations.

Like *Learners*, *Informed Foragers i*) have to decide on a giving-up threshold when they enter a patch, *ii*) within each patch they rely on time between stochastic resource encounters to estimate the local patch quality, and *iii*) they update their giving-up threshold only as they shift between patches. As a consequence, the behavioural mechanisms of *Learning* and *Informed Forgers* differ only in the way they obtain global knowledge. To learn from experience imposes costs that have two major components in our model:

1) *Cost of sampling* relates to lost foraging opportunities during periods when patch-sampling frequency of *Learners* exceeds that of *Informed Foragers*. *Learners* need to sample several patches to collect the information that *Informed Foragers* have for free. As a consequence they spend more time outside foraging patches.

2) *Cost of uncertain estimates* represents the loss in resource accumulation due to time delays and sampling errors when *Learners* estimate giving-up thresholds from experience.

Both costs are emergent properties in our model and they partially depend on the inherited learning rule parameters (γ and τ_0).

Fixed Foragers and the value of learning

A *Fixed Forager* uses a single giving-up threshold throughout the foraging season, and hence ignores information on temporal changes in average resource conditions. It only responds to decreasing encounter frequencies caused by resource depletion in the current patch. We found the best giving-up threshold τ_F with highest average lifetime performance using dynamic programming (detailed in Appendix 1).

Table 1. Variable and parameter definitions. Values and ranges are given for the standard model scenario and values in parenthesis give the range of parameter for which the model was tested.

Variable or parameter	Description	Standard value(s) (range)
t	Time since start of foraging season	
T	Time horizon (total number of time-steps)	2500 (1000-5000)
d	Travel time between patches	25 (5-100)
R_t	Resource level in environment (per patch) at time t	5-55 (0-60)
R_0	Initial resource level in patches	5 (0-60)
r_t	Resource level in local patch	0-55 (0-60)
g	Resource gain per time-step and patch	0.02 (-0.025-0.025)
a	Search efficiency of a forager	0.01
p	Patch number	
P_e	Encounter probability	
k	Total number of resources encountered in a patch	
i	Encounter number in patch	
n_i	Time since last resource encounter	
m	Mortality rate	0.000 (0.000-0.005)
γ	Learning factor	0.00-1.00
τ_0	Initial giving-up threshold of *Learning Foragers*	1.0-50.0
τ_p	Giving-up threshold in patch p for *Learning Foragers*	
τ_F	Giving-up threshold of *Fixed Foragers*	1 - 50
$\tau^*(t)$	Optimal giving-up threshold of *Informed Foragers* at time t	1 - 50
M_m	Mutation rate	0.01
M_s	Mutation step length	0.5-10 %
M_r	Recombination rate	0.1 (0.0-0.5)
f	Reproductive value of one resource encounter	
F	Maximum expected reproductive output using the *Informed* strategy	
V	Reproductive value for *Informed* and *Fixed Foragers*	
$I_{n_i=\tau}$	Patch-leaving indicator function	0 or 1

The fitness difference between *Learning* and *Fixed Foragers* represents the benefit of sampling environmental information and learning from patch experiences. Note that for $\gamma = 0$ the learning strategy is identical to a fixed-rule strategy with $\tau_F = \tau_0$.

For each strategy, we simulated the behaviour of 5000 individuals and compared average giving-up thresholds (strategies), patch-leaving frequencies (foraging behaviour), and resource intake (foraging performance, proportional to reproductive output) at different times within the foraging season. Individual foragers do not interact, meaning that resource levels are unaffected by the resource consumption of other foragers. We found learning factors and giving-up thresholds for different combinations of resource gain rates and mortality regimes (parameter ranges as in Table 1).

Results and discussion

Learners behave differently from both *Fixed* and *Informed Foragers*. In the following, we will 1) show that different strategies result in behavioural changes throughout the season, 2) illustrate the temporal distribution of costs and benefits of learning, and 3) show that life expectancy may affect the value of information and thus change behaviour and foraging efficiency.

Patch allocation strategies and sampling rates

The giving-up threshold of *Fixed Foragers* is constrained to stay constant the whole season. For strategies that utilize various degrees of information, on the other hand, the threshold varies as a response to changing resource levels (Fig. 1a). In an environment where resource conditions become increasingly better throughout the season, there are pronounced differences between the patch-leaving strategies of *Learning* and *Informed Foragers*. Early in season, *Learners* sample the environment to adjust to resource conditions and track temporal resource fluctuations. This sampling activity has

consequences at two levels. First, *Learning Foragers* need to reside sufficiently long in a patch to estimate local patch quality. Second, they need to visit several patches to estimate the global resource conditions. As *Learners* harvest both resources and information, time allocated in each patch represents a trade-off between resource exploitation and information gathering. The genetically determined initial giving-up threshold is therefore relatively low, which ensures that individuals make frequent movements and thereby sample more patches. Under poor environmental conditions, *Learning Foragers* will continue to encounter patches with relatively few resources. Giving-up thresholds then increase and approach that of *Informed Foragers* (Fig. 1a). Later in season giving-up thresholds of *Learners* will commonly lag somewhat behind that of *Informed Foragers*. This is because *Learners* need to make experiences before they update their strategy.

Differences in giving-up thresholds are reflected in patch residence times (Fig. 1b). Early in the foraging season, *Learners* have short patch residence times and change patches more frequently than *Informed Foragers*. After an initial exploration phase, *Learners* generally reside somewhat longer in each patch. When resource conditions improve during the season, giving-up thresholds decrease and *Informed Foragers* leave patches of increasingly higher quality. As a result, patch residence times of *Informed Foragers* change only slightly during the foraging season.

The value of learning

Since reproductive output is proportional to lifetime resource accumulation in our model, differences in resource intake between *Learning* and *Fixed Foragers* represent the value of learning. The relative performance of each strategy changes through the season as the *Learning* and *Fixed Forager* strategies take turns in being closest to the *Informed* strategy (Fig. 1). The value of each strategy therefore has to be averaged over the entire season, discounting for the probability that the forager will die before the potential foraging benefit is realized. As long as the benefits of updating resource estimates outweigh the costs of sampling, there will be a value of acquiring information from patch experiences (Fig. 1c).

All foragers have to choose a giving-up threshold when they enter a new patch. At any time in season, both the populations of *Learning* and *Informed Foragers* will span over individuals with different thresholds: *Learning Foragers* because of differences in individual experiences, and *Informed Foragers* because they entered patches at different times. In contrast, *Fixed Foragers* have a single giving-up threshold, and will be close to the optimal strategy only for a short time period of the season (Fig. 1a). During this period, *Fixed Foragers* have higher mean intake rates than the other strategies since *all* individuals in the *Fixed Forager* population follow the optimal strategy (Fig. 1c). Variation among individuals represents a general cost for flexible
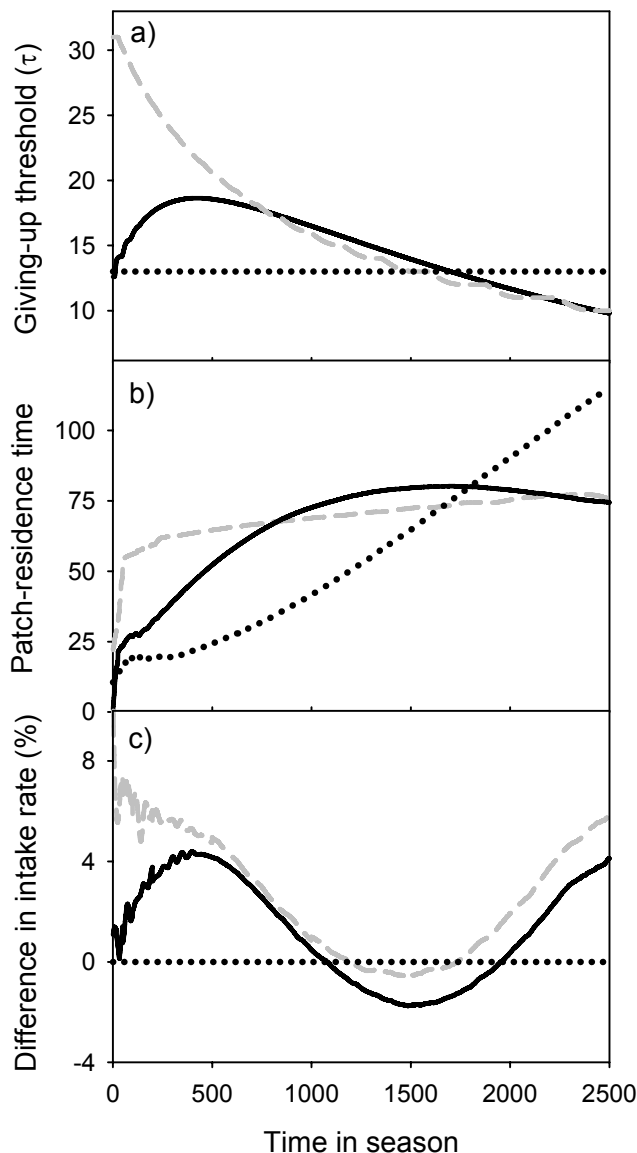
Figure 1. Patch-leaving strategies, patch-allocation behavior and relative foraging performance as a function of time in season. *Fixed Foragers* (black dotted) are restricted to use the same giving-up threshold throughout the season. *Learning Foragers* (black solid) sample the resource environment and update their giving-up threshold based on experiences. *Informed Foragers* (grey hatched) have free information on seasonal changes in resource levels and select the best giving-up threshold upon entering a new patch. a) Average giving-up thresholds τ b) Average patch-residence times. c) Differences in resource intake rates of *Informed* and *Learning Foragers* relative to *Fixed Foragers* change as resource conditions gradually improve through the season ($g = 0.02$, $R_0 = 5$).

strategies, and is most pronounced in *Learners* that acquire information from individual patch quality estimates.

Cost of learning

Learning and *Informed Foragers* have the same uncertainty when they assess the resource level within a patch. *Informed Foragers*, however, have free information about changes in average resource levels – the property that *Learners* aim at estimating. Resource intake rates of *Learners* are therefore constantly below that of *Informed Foragers* (Fig. 1c). This difference represents the cost that *Learning Foragers* pay for not having *a priori* information. First, *Learners* pay a time cost of frequent sampling as they forego foraging opportunities by sampling more patches. This leads to a relatively low intake rate early in the season (Fig. 1c). Later, *Learners* still have to sample the environment, but now learning costs are more related to imperfect information. Due to the stochastic nature of the resource encounter process, foragers make sampling errors. This introduces uncertainty in estimates of giving-up thresholds and results in suboptimal patch residence times.

It is important to note that differences in strategy or patch allocation behaviour do not translate directly into differences in fitness. Extensive sampling behaviour has limited consequences on intake rates early in season, since resource conditions are relatively poor for the progressively improving environment in this case. Towards the end of the season, *Learners* get steadily better at estimating the optimal giving-up threshold, but the divergence in intake rates stays relatively constant since resource conditions continuously improve. It will be advantageous to prepare for this late and prosperous period whenever foragers have long life expectancies.

Foraging strategies change with life expectancy

The costs and benefits of learning are to some extent separated in time. What can be considered an evolutionarily adaptive patch-leaving strategy therefore depends on survival prospects. When mortality rates increase, early resource harvest becomes more important and foragers discount future foraging opportunities. As a consequence, *Informed Foragers* change their strategies towards slightly higher giving-up thresholds (Fig. 2a). For *Learning Foragers*, however, the strategy changes much more. First, the initial learning phase is compressed and patch sampling frequency drops (Fig. 2b).
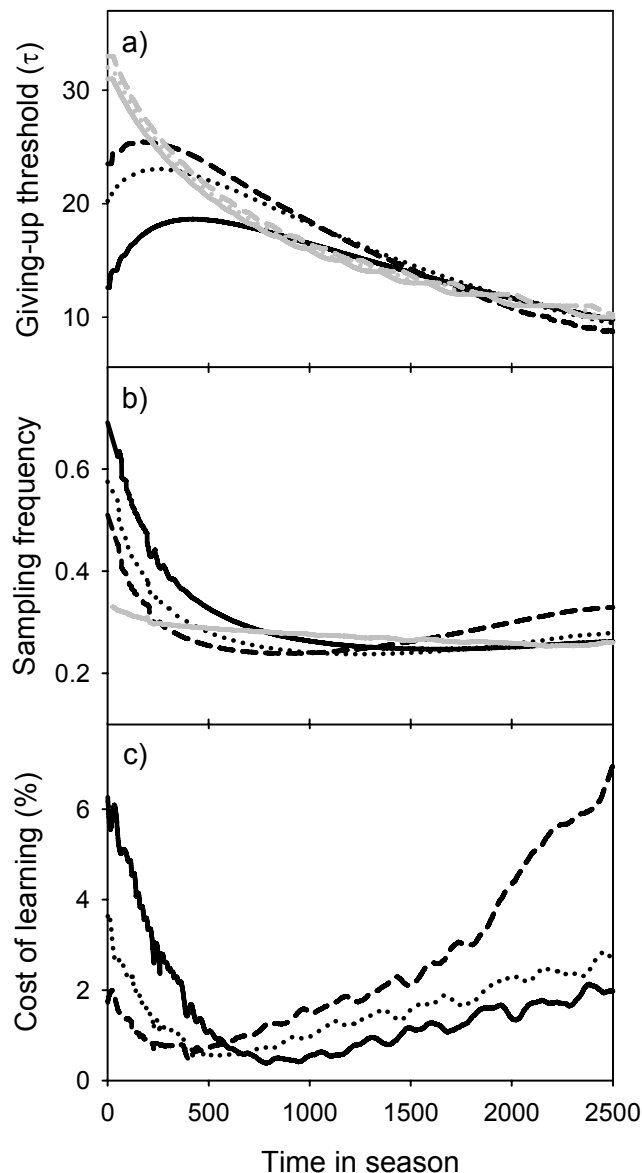

Figure 2. Patch-leaving strategies, sampling frequencies and learning costs as a function of time in season. *Learning* (black lines) and *Informed Foragers* (grey lines) are adapted to high (solid, $m = 0$), intermediate (dotted, $m = 0.001$), and low (hatched, $m = 0.002$) lifetime expectancies. a) Average giving-up thresholds τ in populations of *Learning* and *Informed Foragers*. b) Number of foragers that enter a new patch, relative to the maximum number of patch shifts that foragers travelling for $d = 25$ time steps can potentially make. The relative frequency of patch shifts of *Informed Foragers* is given as a reference (grey, $m = 0.000$). c) Relative differences in resource intake between *Learning* and *Informed Foragers* represent the emergent costs of information acquisition. Costs are paid at different times for foragers adapted to various mortality regimes. During early sampling, learning costs are mainly associated with low patch-residence times, whereas uncertain estimates reduce foraging performance later in season. (Standard parameter values as listed in Table 1).

This is achieved by a higher initial giving-up threshold that more quickly approaches the informed strategy (Fig. 3a). Next, adaptive learning strategies have higher learning factors that bias estimates towards recent experiences. This potentially increases the rate at which *Learners* may adjust their giving-up thresholds (Fig. 3b).

Evolutionarily adaptive learning behaviour depends on the rate of change in environmental conditions (Fig. 3). When resource conditions are relatively stable (gain rates close to zero), foragers evolve fixed strategies with giving-up thresholds similar to the informed strategy. Learning factors generally increase with environmental variation, as has been found also in previous studies (McNamara & Houston, 1985, 1987). The effect is, however, more pronounced in high mortality regimes.

When life-time expectancy decreases, fixed foraging strategies evolve even at moderate seasonal changes in resource levels (Fig. 3b). It is not a shorter life span itself that reduces the value of learning, but rather a shift in the potential costs and benefits of information acquisition (Dukas & Visscher, 1994). In environments where
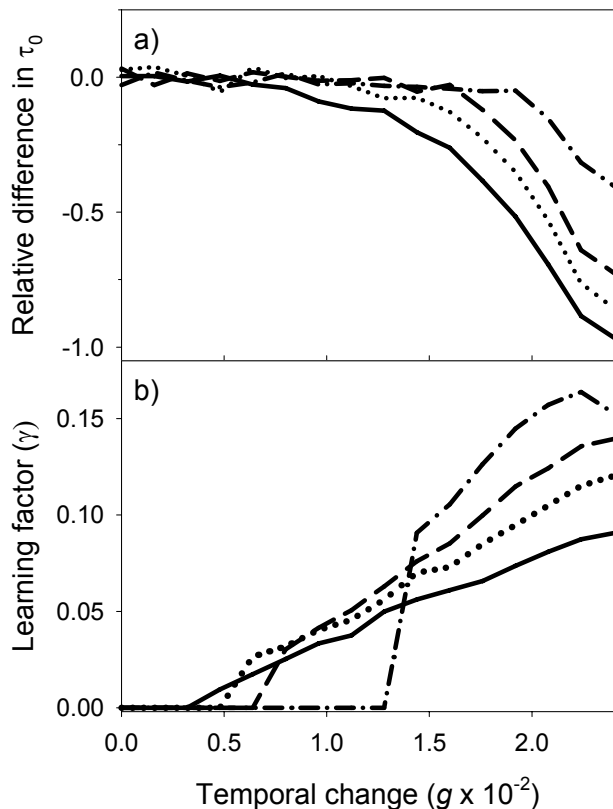

Figure 3. Inherited traits (τ_0 and γ) of *Learning Foragers* adapted to different mortality regimes (solid: $m = 0.000$, dotted: $m = 0.001$, hatched: $m = 0.002$, hatch-dotted: $m = 0.004$). The average resource level equals 30 resources per patch in all environments. Learning rule parameters vary with the rate of temporal change (g) in resource level: a) Differences between the inherited giving-up thresholds τ_0 of *Learning Foragers* and the corresponding initial threshold $\tau^*(1)$ of *Informed Foragers*. b) Learning factors γ of *Learning Foragers* adapted to environments with different life-time expectancies.

resource levels decrease through the season, foragers attain fixed threshold strategies (results not shown). This is because patch sampling then has to take place during the period of prosperous resource conditions early in season, which increases the cost of exploration relative to the benefit that can be attained by enhanced exploitation later on. As long as initial resource conditions do not change between foraging seasons, the model predicted a fixed threshold strategy which is tuned to exploiting the rich resource conditions early in life.

Investment in learning depends on future prospects

Life expectancy affects the behavioural trade-off between resource exploitation and habitat exploration. When mortality increases, *Learners* invest more in present resource consumption at the expense of information harvesting. As a consequence, sampling costs decrease early in season (Fig. 2c). Less sampling, however, leads to less accurate estimation (Fig. 4) and reduces foraging performance of *Learners* later in season (Fig. 2c).
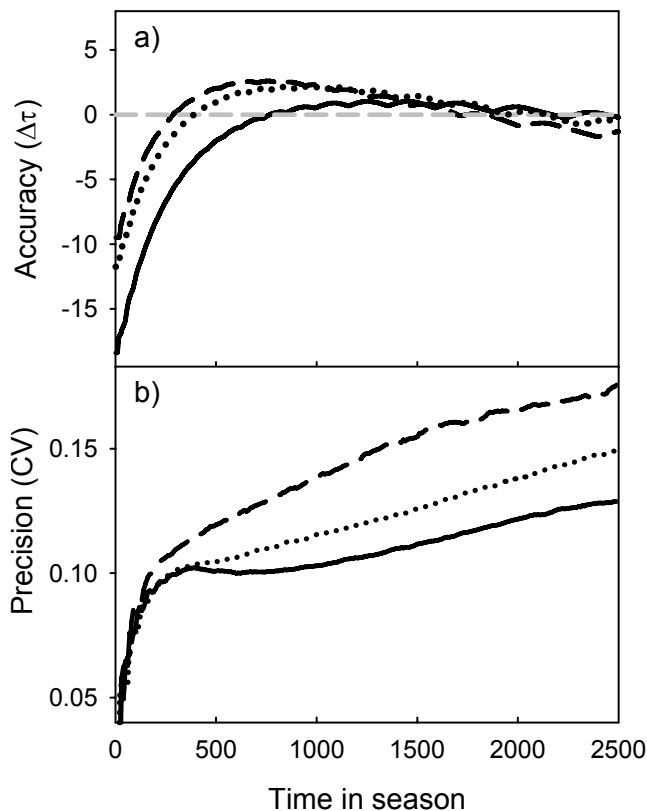
Figure 4. Precision and accuracy in estimates of giving-up thresholds of *Learning Foragers* adapted to different mortality regimes. Solid lines represent foragers adapted to environments with no mortality (m = 0.000, τ = 12.7, γ = 0.078); dotted lines represent intermediate life-time expectancies (m = 0.001, τ = 18.9, γ = 0.096); and dashed lines low survival prospects (m = 0.002, τ = 23.5, γ = 0.129). a) Deviations in giving-up threshold ($\Delta\tau$) between the *Learning* (black lines) and *Informed Foragers* (grey dashed line) as a function of time in the foraging season. b) Coefficient of variation (CV) in estimates of giving-up threshold in a population of 10 000 *Learning Foragers* adapted to different mortality regimes.

Three perspectives can provide a better view of the details and mechanisms that higher mortality incurs. First (motivation): the expected resource intake at any time in the season equals the intake rate achieved by a forager multiplied by the probability of survival up to that time. Consequently, sacrificing present foraging opportunities to improve performance later in season becomes less advantageous as life expectancy decreases. Second (mechanism): *Learners* evolve higher initial giving-up thresholds when mortality rate increases, which reflects how they invest in early foraging. As the initial giving-up threshold of *Learners* approaches that of *Informed Foragers*, the behaviour of the two strategies becomes more similar early in the season (Fig. 4a). Thus, *Learners* spend more time exploiting patches, which results in lower sampling activity and reduced information updating frequencies. To compensate, learning factors increase and consequently each individual patch experience is given more weight. Third (consequence): estimates of giving-up thresholds have lower precision and vary more among *Learners* when risk of mortality is high (Fig. 4b). Shorter lifetime expectancy selects for learning strategies that produce estimates with reduced accuracy and lower precision. This results in relatively high resource intake rates early in the season, when extensive sampling typically incurs high costs on foragers adapted to low mortality environments (Fig. 2c).

General discussion

Learners harvest both resources and information in this model. A period of habitat exploration commonly precedes a more intensive resource exploitation phase. During information harvesting, *Learners* are primarily motivated by the need to gain experience and canalize learning in right directions. Experimental studies of patch allocation and forager distributions have commonly omitted such early exploratory phases from analysis (but see e.g. Krebs *et al.*, 1978; Shettleworth *et al.*, 1988). Early models in behavioural and evolutionary ecology commonly assumed that gains

derived form different strategies alternatives should be evaluated in immediate food intake rates (Emlen, 1966; MacArthur & Pianka, 1966). A generally criticism of such models is that they ignore the benefit of acquiring information that may enhance future foraging or reproductive success.

Learning represents an investment that improves future foraging performance. It affects resource harvest directly, as individuals alter behaviour based on foraging experiences, and indirectly, as accuracy and precision of resource estimates influence behavioural decisions. The way that behaviour of *Learners* deviate from predictions based on *Informed* or other types of omniscient foragers (Nonacs, 2001), may therefore depend on 1) the period in which behaviour is observed, and 2) characteristics of the forager's natural environment. We demonstrate how survival prospects may alter the trade-off between exploration and exploitation. Other ecological factors, such as seasonal variation in offspring survival and risk of starvation, may influence the value of information in similar ways. This illustrates how adaptive foraging strategies optimize rather than maximize the quality of resource estimates, and that the optimization criterion is not information precision but fitness. An additional aspect not addressed here, is the robustness of a strategy in obtaining basic requirements which become increasingly important in unpredictable and variable environments (Mangel, 1990; Inglis *et al.*, 2001; Dall & Johnstone, 2002; Carmel & Ben-Haim, 2005).

Learners need to integrate past and present information to efficiently track temporal change. The time period that information is retained in memory (Devenport & Devenport, 1994; Devenport *et al.*, 1997; Devenport, 1998; Hughes & Blight, 1999; Devenport *et al.*, 2005) and the way recent patch experience influences behavioural choices vary with characteristics of the natural environment (van Baaren *et al.*, 2005). Neural structures involved in memory storage and learning are metabolically costly, hence energy supply may potentially constrain coding and processing of sensory information (Bernays, 1998; Laughlin *et al.*, 1998; Dukas, 1999; Laughlin, 2001). In *Drosophila melanogaster,* larval competitive ability and adult resistance to desiccation

and starvation are lower in individuals with genetically high learning capacity (Mery & Kawecki, 2003, 2005). This illustrates how formation and maintenance of memory records impose ecologically relevant costs.

In parasitoids, patch-leaving behavior ranges from that mainly based on innate rules to that which is mostly learned from experience (Vos *et al.*, 1998; Wajnberg *et al.*, 2000; Boivin *et al.*, 2004). Depending on the spatial distribution of their host, female parasitoids may either adjust their behaviour based on recent patch experiences or use a non-responsive fixed strategy (Vos *et al.*, 1998).

A simple giving-up time rule may arrive at sub-optimal patch leaving decisions in highly stochastic environments (Oaten, 1977; Iwasa *et al.*, 1981; Green, 1984). Hence, foragers that integrate their experiences of resource encounters with additional sensory information could be provided with better patch quality estimates (Iwasa *et al.*, 1981; Green, 1984; Valone, 1989; Persons & Uetz, 1996; van Alphen *et al.*, 2003; Olsson & Brown, 2006). Simple rules of thumb may, however, often work satisfactorily (as Green, 1984 illustrated for the giving-up time rule) and be more biologically realistic as they offer quick responses to environmental cues (Bernays, 1998; Hutchinson & Gigerenzer, 2005). Rules of thumb normally require less computational effort than for instance Bayesian updating rules (Gigerenzer & Todd, 1999). In addition, foragers are generally better at accurately estimating and discriminating between shorter time periods (Shettleworth, 1998; Stephens, 2002).

Under natural conditions, learning abilities are commonly biased towards certain experiences (Arak & Enquist, 1998; Dukas, 2002). Genetic predisposition and prior knowledge are, however, seldom considered in studies of foraging and memory (McNamara *et al.*, 2006). Our results illustrate that ecological factors may influence inherited traits and the way prior expectations are updated based on new information. Instead of incorporating an inherited patch allocation strategy adapted to *forage* efficiently early in life, *Learners* were genetically predisposed *to learn* to behave efficiently. When information acquisition is considered in an ecological context, learning experiences become a function of the behavioural strategy itself. Prior

expectations hence function to canalize information acquisition and not only to give the best possible estimate from the start. This emphasis how animals are more likely to learn from situations that ancestors have been exposed to, simply because they attend to and can interpret relevant information (McNamara *et al.*, 2006).

Acknowledgements

References

Abrahams, M.V. (1986) Patch choice under perceptual constraints: a cause for departures from an ideal free distribution. *Behav Ecol Sociobiol*, **19**, 409-415.

Arak, A. & Enquist, M. (1993) Hidden preferences and the evolution of signals. *Philos Trans R Soc Lond B Biol Sci*, **340**, 207-213.

Arak, A. & Enquist, M. (1998) Neural representation and the evolution of signal form. In *Cognitive ecology: the evolutionary ecology of information processing and decision making* (ed R. Dukas), pp. 21-87. The Univeristy of Chicago Press, Chicago, MI.

Bernays, E.A. (1998) The value of being a resource specialist: Behavioral support for a neural hypothesis. *Am Nat*, **151**, 451-464.

Bernstein, C., Kacelnik, A., & Krebs, J.R. (1988) Individual decisions and the distribution of predators in a patchy environment. *J Anim Ecol*, **57**, 1007-1026.

Boivin, G., Fauvergue, X., & Wajnberg, E. (2004) Optimal patch residence time in egg parasitoids: innate versus learned estimate of patch quality. *Oecologia*, **138**, 640-647.

Carmel, Y. & Ben-Haim, Y. (2005) Info-gap robust-satisficing model of foraging behavior: do foragers optimize or satisfice? *Am Nat*, **166**, 633-641.

Charnov, E.L. (1976) Optimal foraging, the marginal value theorem. *Theor Popul Biol*, **9**, 129-136.

Clark, C.W. & Mangel, M. (1984) Foraging and flocking strategies: Information in an uncertain environment. *Am Nat*, **123**, 626-641.

Clark, C.W. & Mangel, M. (1986) The evolutionary advantages of group foraging. *Theor Popul Biol*, **30**, 35-75.

Clark, C.W. & Mangel, M. (2000) *Dynamic state variable models in ecology.* Oxford University Press, Oxford, UK.

Cohen, D. (1993) The equilibrium distribution of optimal search and sampling effort of foraging animals in patchy environments. *Lect Notes Biomath*, **98**, 173-191.

Dall, S.R.X., Giraldeau, L.A., Olsson, O., McNamara, J.M., & Stephens, D.W. (2005) Information and its use by animals in evolutionary ecology. *Trends Ecol Evol*, **20**, 187-193.

Dall, S.R.X. & Johnstone, R.A. (2002) Managing uncertainty: information and insurance under the risk of starvation. *Philos Trans R Soc Lond B Biol Sci*, **357**, 1519-1526.

Devenport, J.A., Patterson, M.R., & Devenport, L.D. (2005) Dynamic averaging and foraging decisions in horses (*Equus callabus*). *J Comp Psychol*, **119**, 352-358.

Devenport, L., Hill, T., Wilson, M., & Ogden, E. (1997) Tracking and averaging in variable environments: A transition rule. *J Exp Psychol Anim Behav Process*, **23**, 450-460.

Devenport, L.D. (1998) Spontaneous recovery without interference: Why remembering is adaptive. *Anim Learn Behav*, **26**, 172-181.

Devenport, L.D. & Devenport, J.A. (1994) Time-dependent averaging of foraging information in least chipmunks and golden-mantled ground squirrels. *Anim Behav*, **47**, 787-802.

Dukas, R. (1998) Evolutionary ecology of learning. In *Cognitive ecology: the evolutionary ecology of information processing and decision making* (ed R. Dukas). The University of Chicago Press, Chicago, MI.

Dukas, R. (1999) Costs of memory: Ideas and predictions. *J Theor Biol*, **197**, 41-50.

Dukas, R. (2002) Behavioural and ecological consequences of limited attention. *Philos Trans R Soc Lond B Biol Sci*, **357**, 1539-1547.

Dukas, R. & Visscher, P.K. (1994) Lifetime learning by foraging honey-bees. *Anim Behav*, **48**, 1007-1012.

Emlen, J.M. (1966) Role of time and energy in food preference. *Am Nat*, **100**, 611-617.

Fretwell, S.D. & Lucas, H.L.J. (1970) On territorial behaviour and other factors influencing habitat distribution in birds. *Acta Biotheor*, **19**, 16-36.

Gigerenzer, G. & Todd, P.M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart* (eds G. Gigerenzer, P.M. Todd & the ABC research group), pp. 3-34. Oxford University Press, New York.

Gould, J.P. (1974) Risk, stochastic preference, and the value of information. *J. Econ. Theory*, **8**, 64-84.

Green, R.F. (1984) Stopping rules for optimal foragers. *Am Nat*, **123**, 30-43.

Hirvonen, H., Ranta, E., Rita, H., & Peuhkuri, N. (1999) Significance of memory properties in prey choice decisions. *Ecol Model*, **115**, 177-189.

Holland, J.H. (1992) *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence.* MIT Press, Cambridge, MA.

Houston, A.I. & McNamara, J.M. (1999) *Models of adaptive behaviour - An approach based on state.* Cambridge University Press, Cambridge, UK.

Hughes, R.N. & Blight, C.M. (1999) Algorithmic behaviour and spatial memory are used by two intertidal fish species to solve the radial maze. *Anim Behav*, **58**, 601-613.

Hutchinson, J.M.C. & Gigerenzer, G. (2005) Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behav Process*, **69**, 97-124.

Inglis, I.R., Langton, S., Forkman, B., & Lazarus, J. (2001) An information primacy model of exploratory and foraging behaviour. *Anim Behav*, **62**, 543-557.

Iwasa, Y., Higashi, M., & Yamamura, N. (1981) Prey distribution as a factor determining the choice of optimal foraging strategy. *Am Nat*, **117**, 710-723.

Krebs, J.R. & Inman, A.J. (1992) Learning and foraging: Individuals, groups, and populations. *Am Nat*, **140**, S63-S84.

Krebs, J.R., Kacelnik, A., & Taylor, P. (1978) Test of optimal sampling by foraging great tits. *Nature*, **275**, 27-31.

Laughlin, S.B. (2001) Energy as a constraint on the coding and processing of sensory information. *Curr Opin Neurobiol*, **11**, 475-480.

Laughlin, S.B., van Steveninck, R.R.D., & Anderson, J.C. (1998) The metabolic cost of neural information. *Nat Neurosci*, **1**, 36-41.

MacArthur, R.H. & Pianka, E.R. (1966) On optimal use of a patchy environment. *Am Nat*, **100**, 603-&.

Mangel, M. (1990) Dynamic information in uncertain and changing worlds. *J Theor Biol*, **146**, 317-332.

Mangel, M. & Clark, C.W. (1983) Uncertainty, search, and information in fisheries. *J. Cons. int. Explor. Mer*, **41**, 93-103.

McNamara, J.M. (1982) Optimal patch use in a stochastic environment. *Theor Popul Biol*, **21**, 269-288.

McNamara, J.M., Green, R.F., & Olsson, O. (2006) Bayes' theorem and its applications in animal behaviour. *Oikos*, **112**, 243-251.

McNamara, J.M. & Houston, A.I. (1985) Optimal foraging and learning. *J Theor Biol*, **117**, 231-249.

McNamara, J.M. & Houston, A.I. (1987) Memory and the efficient use of information. *J Theor Biol*, **125**, 385-395.

Mery, F. & Kawecki, T.J. (2003) A fitness cost of learning ability in *Drosophila melanogaster*. *Proc R Soc Lond B Biol Sci*, **270**, 2465-2469.

Mery, F. & Kawecki, T.J. (2005) A cost of long-term memory in *Drosophila*. *Science*, **308**, 1148-1148.

Nonacs, P. (2001) State dependent behavior and the Marginal Value Theorem. *Behav Ecol*, **12**, 71-83.

Oaten, A. (1977) Optimal foraging in patches: A case for stochasticity. *Theor Popul Biol*, **12**, 263-285.

Olsson, O. & Brown, J.S. (2006) The foraging benefits of information and the penalty of ignorance. *Oikos*, **112**, 260-273.

Olsson, O. & Holmgren, N.M.A. (1998) The survival-rate-maximizing policy for Bayesian foragers: wait for good news. *Behav Ecol*, **9**, 345-353.

Papaj, D.R. (1993) Automatic behavior and the evolution of instinct: lesson from learning in parasitoids. In *Insect learning: Ecological and evolutionary perspectives* (eds D.R. Papaj & A.C. Lewis), pp. 243-272. Chapman & Hall, New York.

Persons, M.H. & Uetz, G.W. (1996) The influence of sensory information on patch residence time in wolf spiders (*Araneae: Lycosidae*). *Anim Behav*, **51**, 1285-1293.

Rodriguez-Gironés, M.A. & Vásquez, R.A. (1997) Density-dependent patch exploitation and acquisition of environmental information. *Theor Popul Biol*, **52**, 32-42.

Shettleworth, S.J. (1998) *Cognition, evolution and behavior.* Oxford University Press, Oxford, UK.

Shettleworth, S.J., Krebs, J.R., Stepehens, D.W., & Gibbons, J. (1988) Tracking a fluctuating environment: a study of sampling. *Anim Behav*, **36**, 87-105.

Stephens, D.W. (1987) On economically tracking a variable environment. *Theor Popul Biol*, **32**, 15-25.

Stephens, D.W. (1993) Learning and behavioral ecology: Incomplete information and environmental predictability. In *Insect Learning: Ecology and Evolutionary Perspectives.* (eds D.R. Papaj & A.C. Lewis), pp. 195-218. Chapman & Hall, London, UK.

Stephens, D.W. (2002) Discrimination, discounting and impulsivity: a role for an informational constraint. *Philos Trans R Soc Lond B Biol Sci*, **357**, 1527-1537.

Sumida, B.H., Houston, A.I., McNamara, J.M., & Hamilton, W.D. (1990) Genetic algorithms and evolution. *J Theor Biol*, **147**, 59-84.

Valone, T.J. (1989) Group foraging, public information, and patch estimation. *Oikos*, **56**, 357-363.

van Alphen, J.J.M., Bernstein, C., & Driessen, G. (2003) Information acquisition and time allocation in insect parasitoids. *Trends Ecol Evol*, **18**, 81-87.

van Baaren, J., Boivin, G., & Outreman, Y. (2005) Patch exploitation strategy by an egg parasitoid in constant or variable environment. *Ecol Entomol*, **30**, 502-509.

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. *J Anim Ecol*, **67**, 774-783.

Wajnberg, E., Fauvergue, X., & Pons, O. (2000) Patch leaving decision rules and the Marginal Value Theorem: an experimental analysis and a simulation model. *Behav Ecol*, **11**, 577-586.

Appendix 1

Informed and Fixed foraging strategies

A forager's contribution to the next generation depends on resource intake rate and survival prospects at different times in a foraging season. Foragers continuously reproduce proportionally to their accumulated resources, and we assume that the value

of a resource in terms of reproductive output is *f*. Foragers have two behavioural options; they may either stay in a patch and search for resources, or leave and travel to a new patch. In the same way as *Learners*, *Informed* and *Fixed Foragers* will leave a patch when time since last resource encounter n_i exceeds a giving-up threshold τ. Strategies only differ in the way they obtain their giving-up thresholds.

Informed foragers have free global information and know the best giving-up threshold to use when they enter a new patch. Foragers update this information only as they shift between patches, in much the same way as *Learning Foragers* update their estimates of τ_p upon leaving a patch. We seek the giving-up threshold that maximizes expected lifetime reproductive output for a forager that enterers a patch at a time *t* in a foraging season of length *T*:

$$F(t) = \max_{\tau} E\{\text{accumulated reproduction from } t \text{ to } T\}$$

The expected reproductive output is maximized over all possible giving-up times τ that a forager may select when it enters a patch at time *t*. Foragers are restricted to use the same threshold during the entire patch visit, but may change giving-up threshold when it enters a new foraging patch.

The value of staying in the patch depends on the expected fitness consequences of *i*) encountering a resource in the next time step, and *ii*) the value of resource encounters from time *t*+1 until the end of season. The probability P_e that a forager will encounter a resource in the next time step will depend on the global resource level R_t and the number of resources consumed *k* so far in the patch ($r_t = R_t - k$). For a given time *t* in the season, the value of utilizing strategy τ, V_τ, is therefore given by:

$$V_\tau(t,k,n_i) = e^{-md} F(t+d) I_{n_i=\tau}$$
$$+ \left[P_e(r_t)\{f + e^{-m}V_\tau(t+1,k+1,1)\} + (1-P_e(r_t))e^{-m}V_\tau(t+1,k,n_i+1) \right](1 - I_{n_i=\tau})$$

For a given strategy τ, the indicator factor $I_{n=\tau}$ equals 1 if $n_i \geq \tau$ and is 0 otherwise. The first part represents the value of leaving the patch. The probability that the forager

survives to enter a patch after d time steps of travelling is e^{-md}, where m is mortality rate. The *Informed Forager* selects the optimal giving-up threshold when it enters a new patch, hence the expected future reproductive value from the time it enters the new patch and onwards is $F(t+d)$. The second part of the equation gives the value of staying, which is the sum of *i*) the probability that the forager encounters a resource when it searches a patch multiplied with the expected future reproductive value given that a new resource is consumed, and *ii*) the corresponding value if no resources are encountered in time-step t.

At the time horizon T the reproductive output is $F(T) = 0$. As long as $t+d > T$ a forager will not reach a new patch before the end of season, and the best option is to stay in the patch. Earlier in season, there is an optimal giving-up threshold $\tau^*(t)$ that maximizes future reproductive value for a forager that *enters* a patch at a given time-step t:

$$F(t) = \max_{\tau}\{V_{\tau}(t,0,1)\}$$

Starting at the time horizon T and working backwards, we may calculate the optimal giving-up threshold $\tau^*(t)$ for all time-steps in the season (Clark & Mangel, 2000).

Fixed Foragers use a single, genetically determined, giving-up threshold during the whole foraging season. At any time t in the season the expected reproductive output is:

$$V_{\tau_F}(t,k,n_i) = e^{-md}V_{\tau_F}(t+d,0,1)I_{n_i=\tau_F}$$
$$+ \left[P_e(r_t)\left\{f + e^{-m}V_{\tau_F}(t+1,k+1,1)\right\} + (1-P_e(r_t))e^{-m}V_{\tau_F}(t+1,k,n_i+1)\right](1-I_{n_i=\tau_F})$$

As for *Informed Forages*, the first part gives the value of leaving a patch, and the second the value of staying. Note that the value of leaving is conditional on a fixed giving-up threshold, as opposed to *Informed Forgers* that select the optimal giving-up threshold every time they enter a new foraging patch. We search for the fixed τ_F that had the best average performance and that maximized expected lifetime reproductive output. Starting from the time horizon T, where future reproductive value is zero, we can calculate the value of using different strategies at all times t in a season. We can

hence find the τ_F that maximizes lifetime reproductive output over a season, that is:

$$\max_{\tau_F}\left\{V_{\tau_F}(0,0,1)\right\}.$$