

Evolution of a flexible rule for foraging that copes with environmental variation

Andrew D. HIGGINSON^{*}, Tim W. FAWCETT, Alasdair I. HOUSTON

School of Biological Sciences, Life Sciences Building, University of Bristol, 24 Tyndall Avenue, Bristol, BS8 1TQ, U.K.

Abstract Models of adaptive behaviour typically assume that animals behave as though they have highly complex, detailed strategies for making decisions. In reality, selection favours the optimal balance between the costs and benefits of complexity. Here we investigate this trade-off for an animal that has to decide whether or not to forage for food – and so how much energy reserves to store – depending on the food availability in its environment. We evolve a decision rule that controls the target reserve level for different ranges of food availability, but where increasing complexity is costly in that metabolic rate increases with the sensitivity of the rule. The evolved rule tends to be much less complex than the optimal strategy but performs almost as well, while being less costly to implement. It achieves this by being highly sensitive to changing food availability at low food abundance – where it provides a close fit to the optimal strategy – but insensitive when food is plentiful. When food availability is high, the target reserve level that evolves is much higher than under the optimal strategy, which has implications for our understanding of obesity. Our work highlights the important principle of generalisability of simple decision-making mechanisms, which enables animals to respond reasonably well to conditions not directly experienced by themselves or their ancestors [*Current Zoology* 61 (2): 303–312, 2015].

Keywords Behavioural mechanisms, Fat storage, Fluctuating environments, Generalisation, Obesity, Rule of thumb

Evolutionary models of animal behaviour have typically assumed that behavioural strategies can be highly flexible and detailed (McNamara and Houston, 2009). For example, in models of state-dependent behaviour, the optimal strategy found by dynamic programming takes the form of a large look-up table, which specifies what the animal should do for all values of the state variables (Houston and McNamara, 1999). Such approaches therefore implicitly assume that the animal has sufficient flexibility to respond optimally to every possible set of environmental parameters and internal states (Houston and McNamara, 1999). However, it is important to note that dynamic programming is a computational tool for finding the optimal strategy, not a representation of the process by which animals make decisions (McNamara and Houston, 2009). Such an optimal strategy identifies the upper limit on performance, but ignores any potential costs of the decision-making mechanism. The high precision implied by dynamic-programming solutions would surely be costly if implemented by real organisms, in terms of the required sensory apparatus or cognitive processing (DeWitt et al., 1998). Therefore, the strategies used by real animals are likely to be much less precise.

Furthermore, the complex solutions found by dynamic programming are specific to the chosen parameter values and therefore to a certain set of environmental conditions; they do not reveal what the animal should do if conditions change (McNamara and Houston, 2009). To allow flexibility across conditions it is possible to calculate a dynamic-programming solution for each set of conditions; but these solutions are independent, in that there is no generalisation across similar sets of conditions. From an adaptive perspective, this lack of generalisation seems inefficient. Instead, animals are likely to have evolved simple mechanisms that are less cognitively demanding but still perform reasonably most of the time, across a broad range of conditions. Such simple and inexpensive rules ('rules of thumb') could allow animals to generalise across similar conditions, with the cost that performance is not exactly optimal under any conditions. Progress in behavioural research depends on understanding how natural selection balances the costs and benefits of complexity in decision-making rules.

One area in which animal decision-making has been particularly well studied is how animals decide when to eat and how much energy to store in their body (Brodin,

Received Oct. 31, 2014; accepted Jan. 23, 2015.

^{*} Corresponding author. E-mail: a.d.higginson@bristol.ac.uk

© 2015 *Current Zoology*

2007). Theory predicts that animals should use energetic reserves in order to trade off the risks of two sources of mortality: starvation and predation (Lima 1986; McNamara, 1990; McNamara and Houston, 1990; Houston and McNamara 1993). This prediction is consistent with how food availability and predation risk covary with patterns of fat storage in a range of animals (Gosler et al., 1995; Krams et al., 2010; Meijer et al., 1994). In principle, animals could have an optimal rule that precisely specifies when to eat for every possible level of food availability. However, such a complex rule is likely to be costly to implement. Simpler rules that are less sensitive to changes in food availability may do sufficiently well without incurring high costs. Work is required to elucidate how such simple rules may have been designed by natural selection.

Here, we investigate what kinds of foraging rules are favoured by natural selection when the flexibility of the optimal solution is not cost-free, such that costs constrain the evolution of the foraging rule and its resulting behaviour. We evolve a decision rule under varying degrees of cost and compare the outcome to the optimal strategy found by dynamic programming (Houston and McNamara, 1999; Clark and Mangel, 2000). Our aim is to highlight some general principles about mechanistic constraints on decision-making, but as an illustrative case we focus on the overwinter survival of a small bird that has to decide whether to forage or rest. If there is no day/night cycle, the optimal strategy is a threshold value of energetic reserves above which the animal rests and below which it forages. The example of foraging or resting to maximise overwinter survival is used for illustrative purposes because it is the simplest case of energy management we can consider. We assume static conditions (cf. fluctuating conditions, e.g. Higginson et al., 2012), no energetic cost to reserves (cf. linear or accelerating costs, e.g. Lima, 1986; McNamara and Houston, 1990; Witter and Cuthill, 1993), all-or-none foraging (cf. a continuum of foraging intensity, e.g. Houston et al., 1993), no reproductive attempts (cf. one or more reproductive attempts, e.g. McNamara et al., 1991) and no competition with conspecifics (cf. dynamic game models, e.g. Houston and McNamara 1988). This problem has been studied extensively from an optimality perspective and is well understood (e.g. Brodin, 2007). Furthermore, the survival circuit concept (LeDoux, 2012) implies that the problem of finding food can be meaningfully studied in isolation.

We find that a decision rule subject to a cost of flexibility may evolve to be fairly crude, but by virtue of its

simplicity it incurs only a small cost. By generalising across similar sets of conditions the evolved rule performs reasonably well in comparison to the optimal strategy, suffering only a small reduction in survival unless conditions are very challenging. Furthermore, when food is abundant the evolved decision rule causes the animal to store far more fat than would be optimal, which has implications for our understanding of obesity in humans. We discuss the insights from this approach for the future development of models of behaviour based on decision-making mechanisms that animals may actually utilise.

1 Materials and Methods

1.1 The model

We describe the individual-based model following the ODD standard protocol (Grimm et al., 2006).

1.2 Purpose

The purpose of the model is to assess the effect of including costs in the design of a mechanism controlling energy storage in response to food availability, and comparing a simpler mechanism to the fully flexible optimal strategy. Specifically, we are interested in the relationship between the probability of finding a food item in each time step spent foraging, γ , and the threshold level of reserves below which the animal forages and above which it rests.

1.3 State variables and scales

Each individual has a level of energetic reserves, x , where $x = 0, 1, 2, \dots, x_{\max}$, and this can change within its lifetime. Each individual has a set of 20 genetic traits (hereafter 'genes') that each control the threshold x^* over a particular range of γ : gene 1 controls behaviour for $0.25 \leq \gamma < 0.275$, gene 2 controls behaviour for $0.275 \leq \gamma < 0.3$, and so on. Each of these genes can specify any value of x^* between 0 and x_{\max} ; thus the resulting relationship between x^* and γ can be non-linear and non-monotonic. Note that we use the term 'genes' as a clear way to describe how the rule is stored; we are not suggesting that this is how the strategy is coded in real animals. During the course of evolution, we allow the genes to change value and become active or inactive. If a gene is inactive, the individual is assumed to be insensitive to variation in γ at that point; thus it adopts the same forage/rest threshold as for the first active gene immediately below it. (The reverse case, where the threshold for an inactive gene takes same value as the first active gene immediately *above* it, gives the same results). We do not allow the first gene (or equivalently, in the reverse case, the last gene) to be switched off.

Each individual in each generation has a given probability of finding food in each time step, γ . The value of γ is drawn independently at the start of each generation and determines the food availability experienced by all members of the population in that generation. γ is a uniform random variable with minimum γ_{\min} and maximum γ_{\max} . We assume that γ is constant during an individual's lifetime and that individuals are sensitive to the value of γ , adopting the optimal forage/rest decision corresponding to that level of food availability. However, γ has been variable during evolutionary history, so individuals have the capacity to respond appropriately to other values of γ (which their descendants may subsequently encounter). Such a scenario may occur in annual species where the value of γ fluctuates randomly from year to year (for a related problem, see Collins et al., 2006).

1.3 Process overview and scheduling

The model has two nested loops. Each individual has a lifespan of 2,000 time steps. After all individuals have either died or survived for 2000 time steps, the next generation is created. The animal uses c units of energy in metabolism every time step, and if its reserves reach zero the animal is assumed to have died of starvation. If the animal forages in a given time step, with probability γ it finds a food item that contains r units of energy. If foraging, there is a risk of mortality from predation that depends on energetic reserves. Individuals decide to forage or rest depending on their current reserves and the food availability γ , according to the decision rule specified by their genotype.

1.4 Design concepts

Emergence. An individual's strategy is encoded by their set of genes, whose values are completely free to evolve. The only constraint placed on the strategy is that we assume a cost: the incurred cost of flexibility is proportional to the total number of genes that are active (g), such that the total metabolic cost paid by individual i per time step is $p_i = c + \kappa \frac{g_i}{G}$, where G is the total number of genes and κ is the cost of maximum flexibility (i.e. if all genes are active). Therefore, the number of active genes is a proxy for flexibility.

Sensing. All individuals are assumed to have perfect knowledge of the values of the ecological variables (e.g. γ) and their current level of energetic reserves (x).

Interactions. The 2000 individuals do not interact with each other, either directly or indirectly (e.g. by consuming food).

Stochasticity. Finding food is a stochastic process. A

food item is found in a given time step with probability γ if the individual chooses to forage.

Adaptation. Over generations, the strategy encoded by the genes in the population leads to higher survival.

Fitness. Each generation lasts for 2000 time steps. For each individual, we record the time step at which they die (of starvation or predation), or if they survive for the full 2000 time steps.

1.5 Initialisation

In the first generation, we initialise genetic values in one of two ways. In some runs all individuals have only one active gene, with value $x_{\max}/2$. In other runs all genes are active and set to the optimal values found by dynamic programming. Each generation all individuals are initialised (at time $t = 0$) with reserve level $x = 50$.

1.6 Input

Other aspects of the environment are constant across all our analyses: maximum reserves ($x_{\max} = 200$), size of food items ($r = 4$), intensity of predation ($a = 0.001$) and the baseline metabolic energy usage ($c = 1$).

1.7 Submodels

Predation. For individual i with current reserves x_i , the probability of being killed by a predator while foraging, μ_i , is

$$\mu_i = a \left(1 + \frac{x_i}{x_{\max}} \right) \quad (1)$$

where a is a parameter reflecting the intensity of predation. Equation (1) reflects our assumption that heavier animals (larger x) are more vulnerable to predation while foraging (Lee et al., 1996; Witter et al., 1994). This seems reasonable for most birds and could easily be modified to reflect different patterns of mass-dependent predation in other taxa (e.g. fish). We assume that while resting the animal is safe from predators.

Energy. If the animal is not killed by a predator, its reserves change from one time step t to the next according to

$$x_i(t+1) = x_i(t) - c - \kappa \frac{g_i}{G} + r \quad (3a)$$

if it forages and finds food, and

$$x_i(t+1) = x_i(t) - c - \kappa \frac{g_i}{G} \quad (3b)$$

otherwise, subject to $0 \leq x_i \leq x_{\max}$. Note that there is no additional energetic cost of reserves or foraging activity.

Reproduction. We evolve the genetic values X_1, X_2, \dots, X_{20} by standard methods of asexual reproduction with mutation, assuming non-overlapping generations (Hamblin 2013). For each individual in the next generation we randomly select two potential parents (with re-

placement); of these two candidates, the one that survived the longest produces the offspring (or if the lifespans were equal, a parent was chosen at random). This is repeated until all 2000 offspring have been created. Each offspring inherits the 20 genes of its parent, with the possibility of mutation. With a fixed probability m_x per gene per generation, the genetic value X_j for gene j ($j = 1, 2, \dots, 20$) increases or decreases (with equal probability) by one unit, subject to $0 \leq X_j \leq x_{\max}$. With a further fixed probability m_g each gene may become inactive if it is currently active, or active if it is currently inactive. For a gene j that switches to being active, the corresponding foraging threshold X_j initially takes the same value as the first active gene below it and thus the strategy initially remains insensitive to changes in food availability over the range influenced by the gene, until a mutation occurs that changes X_j . We iterated this process for one million generations, by which time the populations tended to be almost monomorphic in the number of active genes (g) and the values of most genes.

1.8 Optimal strategy

In finding the optimal strategy we consider an indefinite series of discrete time steps and assume that the animal is adapted to minimise the long-term sum of two sources of mortality: starvation and predation. The optimal strategy is a series of look-up tables, one for each possible value of γ , that tell the animal whether to forage or rest given its current reserves x . For comparison to the evolved mechanism (i.e. the genetically encoded strategy, explained above) we find the optimal decision (forage or rest) at integer values of reserves x ($0 \leq x \leq x_{\max}$) for 21 evenly spaced values of γ between 0.25 and 0.75 (i.e. 0.25, 0.275, ..., 0.725, 0.75), using standard dynamic-programming methods (Clark and Mangel, 2000; Houston and McNamara, 1999). The individual's reserves change from one time step t to the next according to

$$x_i(t+1) = x_i(t) - c + r \quad (2a)$$

if it forages and finds food, and

$$x_i(t+1) = x_i(t) - c \quad (2b)$$

otherwise, subject to $0 \leq x_i \leq x_{\max}$. This process gives a threshold value of reserves for each γ , $x^*(\gamma)$, below which the animal forages and above which it rests (cf. Houston and McNamara 1993).

1.9 Simulation experiments

To assess the performance of the evolved rule (after one million generations) we simulated a population of 2,000 individuals following that rule and tracked their

survival for 2,000 time steps. We initialised all individuals with reserve level $x = 50$. Simulations ran for each of 21 values of γ evenly spaced between γ_{\min} and γ_{\max} . We assessed how the cost of flexibility affects survival by simulating four alternative scenarios using different cost functions:

(a) *Cost of evolved flexibility.* A metabolic cost is paid for every gene that is active under the evolved strategy, whereas inactive genes are cost-free ($p_i = c + \kappa \frac{g_i}{G}$).

(b) *Cost of full flexibility.* A metabolic cost is paid for every gene regardless of whether it is active or inactive, i.e. individuals pay a cost equivalent to having all 20 genes active ($p_i = c + \kappa$).

(c) *Cost of unconstrained flexibility.* A metabolic cost is paid for each gene that would be active if the strategy were not subject to any costs during its evolution. Preliminary analysis showed that under cost-free conditions the strategy has on average 12 genes active, so we take this cost to be $p_i = c + \frac{12}{20} \kappa$.

(d) *Cost-free flexibility.* No metabolic cost is paid for active or inactive genes, i.e. flexibility is cost-free ($p_i = c$).

For each of these scenarios we ran 10 replicate simulations. For each replicate we recorded the number of survivors after 2,000 time steps, and used this to investigate the effects of κ , γ , the cost of the inaccuracy of the evolved rule, and the cost of flexibility that is actually paid by comparing the above scenarios to one another and to the optimal strategy found by dynamic programming. A summary of the simulations we ran is shown in Table 1.

2 Results

The number of genes and their values converged across the replicate simulations, and populations were generally as close to monomorphic in their decision rules as the mutation–selection balance allowed. Therefore we present in each case the mean rule used by each population. This represents the mean values of X_j when switched on in >99% of the population. Whether we initialised the population with a single active gene or with 20 active genes at their optimal values, or whether the genes affected values of γ above or below them, had no significant impact on the number of active genes that evolved (Fig. 1) nor the values of those genes.

When there was no cost of active genes ($\kappa = 0$) the

evolved rule still had an average of eight inactive genes. This occurred because the 12 active genes approximate

the optimal strategy reasonably closely (Fig. 2A) and there is very little selective pressure on the inactive

Table 1 Summary of the strategy-finding procedure and simulations

Level	Number of values or replicates	Values
Cost of flexibility	13	0.0, 0.001, 0.005, 0.01, 0.02, 0.03, 0.04, 0.05, 0.07, 0.1, 0.2, 0.3, 0.5.
Strategy replicates	10	-
Food availability γ	21	0.25, 0.275, ... 0.725, 0.75
Simulated flexibility cost	4	Cost paid for: zero genes (cost-free flexibility), 12 genes (unconstrained flexibility), 20 genes (full flexibility) or g_i genes (evolved flexibility)
Simulation replicates	10	-
Total simulations	109,200	-

Each level is part of a hierarchy such that we ran a total of 109,200 simulations of evolved populations.

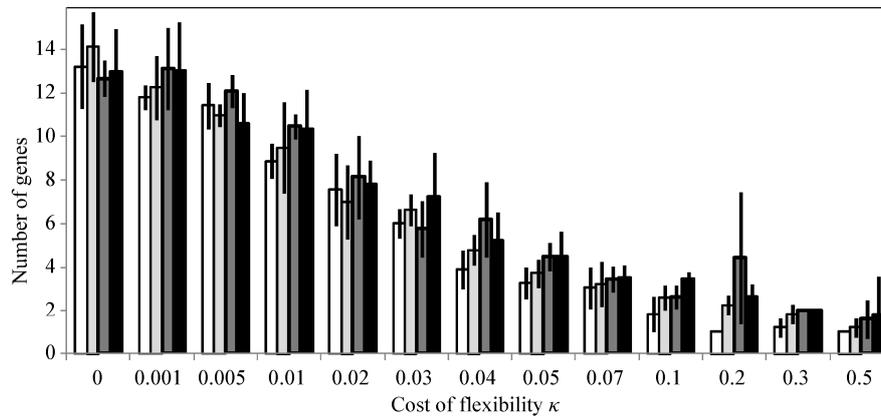


Fig. 1 Mean (± 1 SD) number of active genes as a function of the cost of flexibility κ

The bars are shaded according to whether the algorithm was initialised with just one active gene (white and dark grey bars) or with all 20 active genes of the optimal strategy (light grey and black bars) and whether the threshold for inactive genes matched that of the first active gene above it (dark grey and black bars) or below it (white and light grey bars).

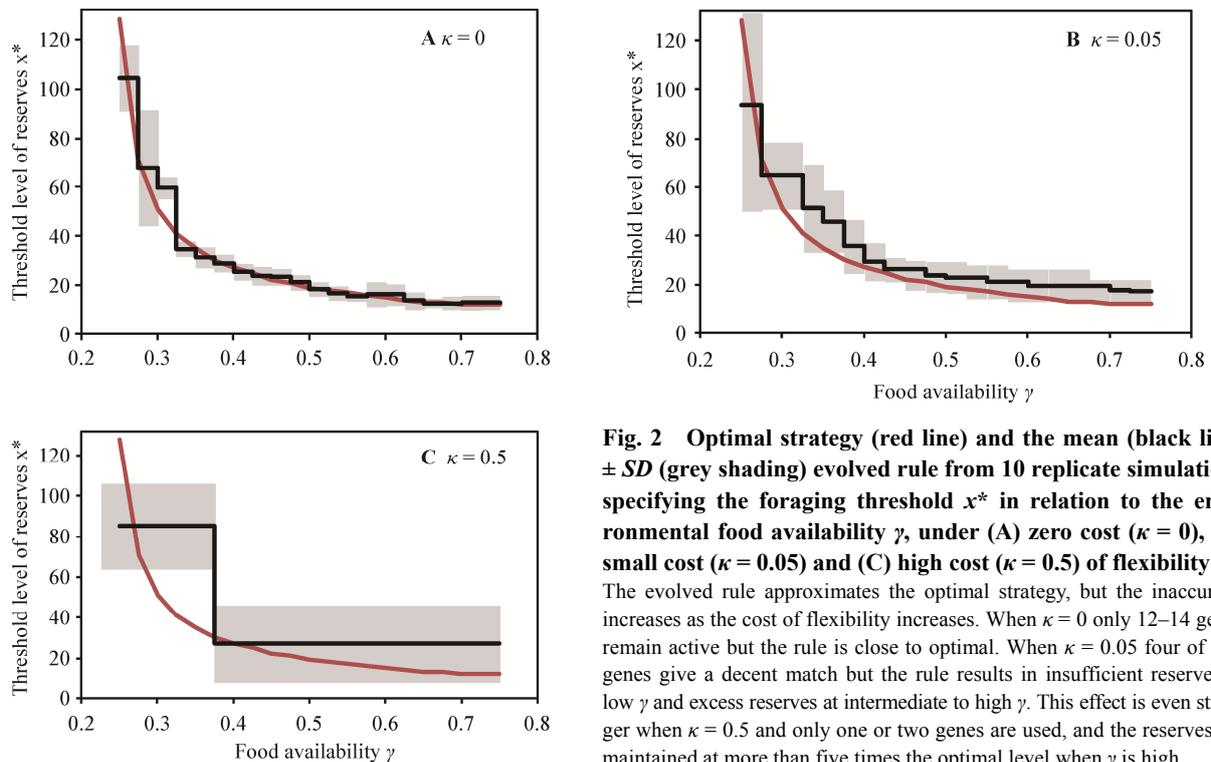


Fig. 2 Optimal strategy (red line) and the mean (\pm SD) evolved rule from 10 replicate simulations specifying the foraging threshold x^* in relation to the environmental food availability γ , under (A) zero cost ($\kappa = 0$), (B) small cost ($\kappa = 0.05$) and (C) high cost ($\kappa = 0.5$) of flexibility

The evolved rule approximates the optimal strategy, but the inaccuracy increases as the cost of flexibility increases. When $\kappa = 0$ only 12–14 genes remain active but the rule is close to optimal. When $\kappa = 0.05$ four of five genes give a decent match but the rule results in insufficient reserves at low γ and excess reserves at intermediate to high γ . This effect is even stronger when $\kappa = 0.5$ and only one or two genes are used, and the reserves are maintained at more than five times the optimal level when γ is high.

genes, because their loss by mutation has a negligible impact on survival (whereas their reactivation is initially selectively neutral; see Methods). However, populations did vary in which genes were active (Fig. 3A), because similar rules can be achieved in many different ways when the majority of genes are active. As expected, an increase in the cost of flexibility, κ , causes a decrease in the number of active genes (Fig. 1).

The average evolved rule when $\kappa = 0.05$ shows clear deviations from the optimal strategy (Fig. 2B), because individuals have on average only four active genes (Fig. 1), mostly concentrated where food availability γ is low (Fig. 3B). The active genes have values close to the optimal strategy, which therefore means that at higher values of γ the animal attempts to carry more reserves than is optimal. When $\kappa = 0.5$ this tendency is even clearer: the evolved rule has just one or two active genes (Fig. 3C) and the animal attempts to carry reserves fivefold larger than is optimal at high γ (Fig. 2C).

We assessed survivorship for the evolved strategy under the four implementations of the incurred cost of flexibility p (see above). Survival decreases with κ at an accelerating rate if the cost of full flexibility (equivalent to 20 active genes) or unconstrained flexibility (equivalent to 12 active genes) is paid; when $\kappa = 0.1$ around 80% of individuals die within 2,000 time steps when $\gamma = 0.3$ (Fig. 4A). By contrast, when individuals only pay a

cost for their evolved level of flexibility, their survivorship is almost as high as when flexibility is cost-free. These patterns are similar at higher values of γ , when more individuals survive (Fig. 4B, C). The cost-free and evolved flexibility conditions are almost as successful as the optimal strategy when $\gamma > 0.3$ (Fig. 5A). That is, although the evolved rule is a coarse approximation of the optimal strategy at high food availability and individuals store far too much energy, survival is not greatly reduced. The cost of an inaccurate mechanism is greater for higher κ (i.e. fewer active genes), but the pattern is otherwise similar (Fig. 5B). Note that relative survivorship under the evolved rule is influenced by which genes tend to be active, as shown by the decrease and then increase in Fig. 4A as γ increases from 0.3.

3 Discussion

Animals do not have unlimited flexibility for responding to variation in their environment, but instead are likely to use simple rules for making decisions (Hutchinson and Gigerenzer, 2005; McNamara and Houston, 2009; Fawcett et al., 2013). A critical aspect of the environment that animals must respond to is the availability of food, and there is a long history of models of adaptive behaviour that predict how animals should use stored energy reserves to cope with stochasticity in finding food (Houston and McNamara, 1999; Brodin,

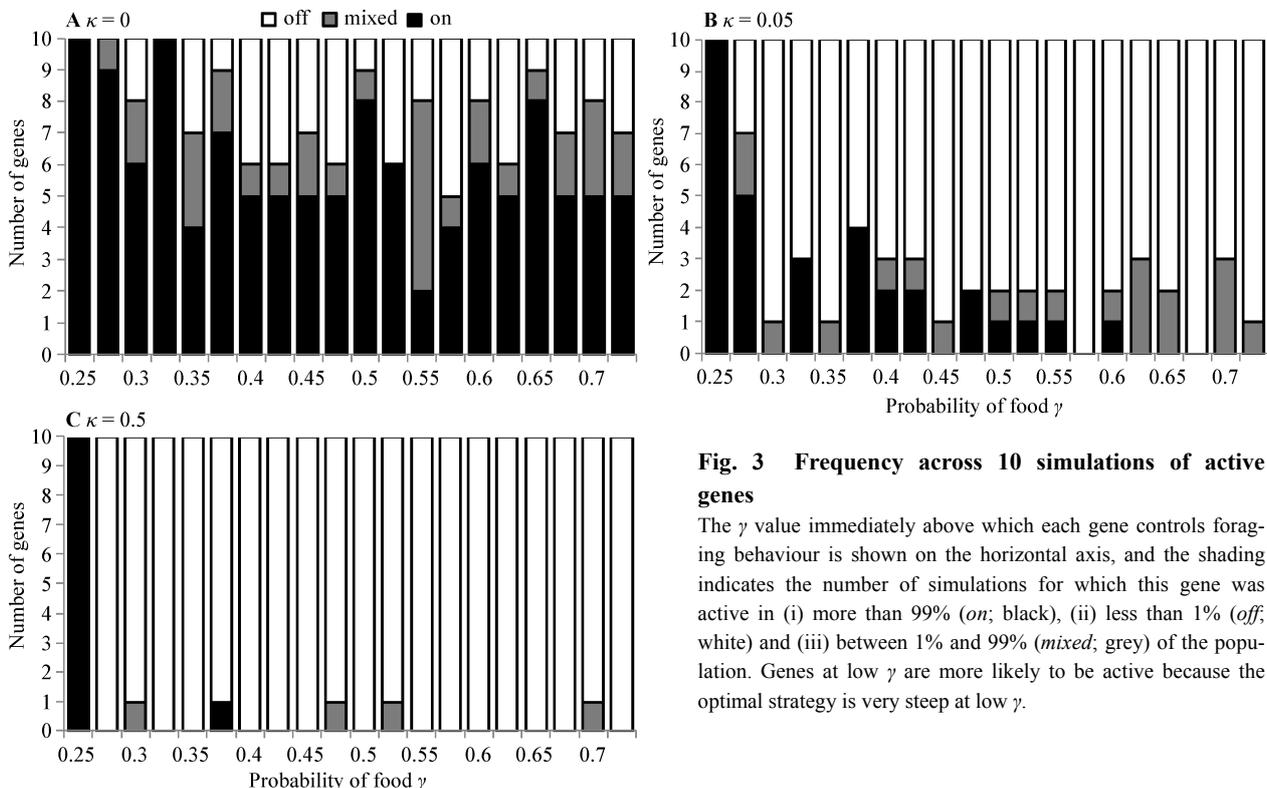


Fig. 3 Frequency across 10 simulations of active genes

The γ value immediately above which each gene controls foraging behaviour is shown on the horizontal axis, and the shading indicates the number of simulations for which this gene was active in (i) more than 99% (*on*; black), (ii) less than 1% (*off*; white) and (iii) between 1% and 99% (*mixed*; grey) of the population. Genes at low γ are more likely to be active because the optimal strategy is very steep at low γ .

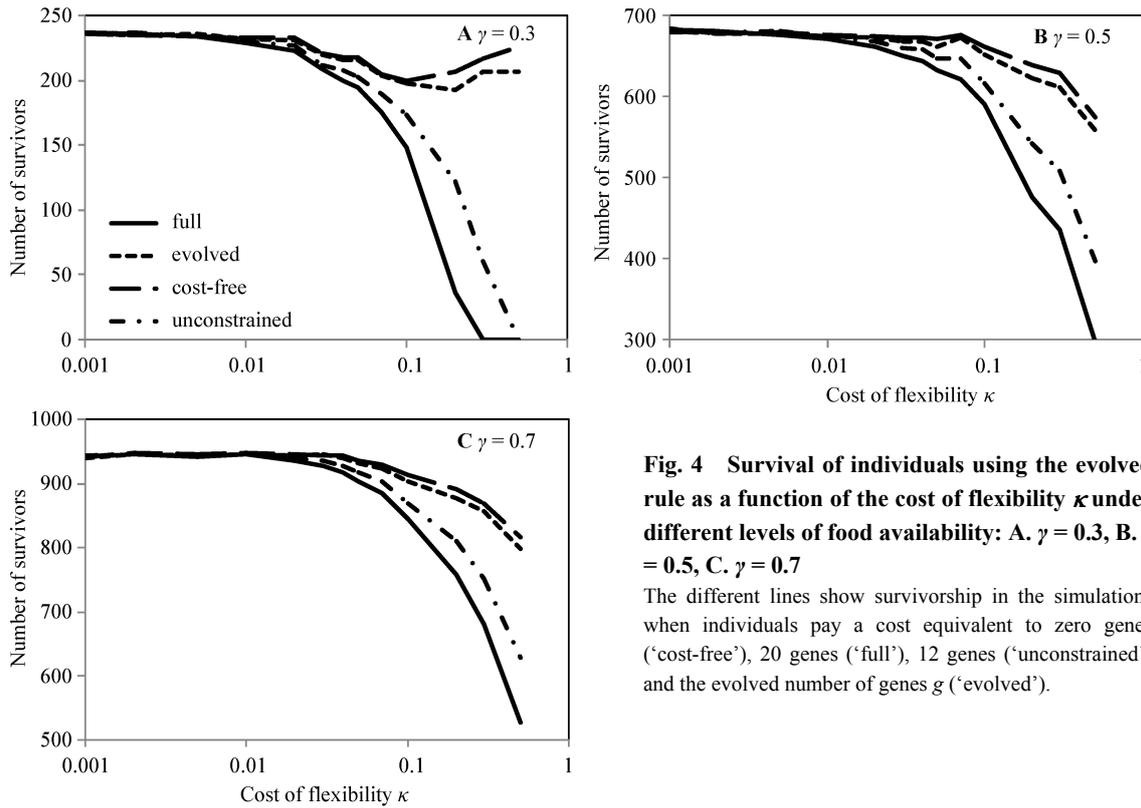


Fig. 4 Survival of individuals using the evolved rule as a function of the cost of flexibility κ under different levels of food availability: A. $\gamma = 0.3$, B. $\gamma = 0.5$, C. $\gamma = 0.7$

The different lines show survivorship in the simulations when individuals pay a cost equivalent to zero genes ('cost-free'), 20 genes ('full'), 12 genes ('unconstrained') and the evolved number of genes g ('evolved').

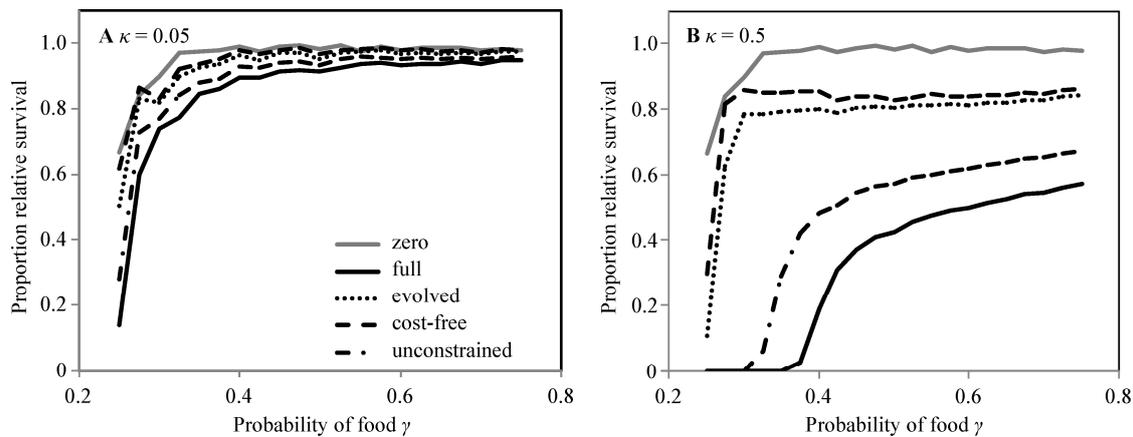


Fig. 5 Survival as a proportion of the number of survivors following the optimal strategy in relation to the environmental food availability γ , for (A) small ($\kappa = 0.05$) or (B) large ($\kappa = 0.5$) costs of flexibility (for comparison, the grey lines show zero cost of flexibility).

The evolved mechanism is not very costly (cf. 'cost-free') and, despite the inaccuracy of the mechanism at high γ resulting in excess reserves, survival is not greatly reduced.

2007). We have evolved a rule for controlling behaviour in a simple model of energy reserve use, where the objective is to minimise the long-term rate of mortality. We find that the evolved rule is much less responsive than the optimal strategy to small changes in food availability. Even when we assumed no cost to flexibility, limited flexibility (12 genes) evolves. This is sufficient to make the evolved rule very close to the optimal strategy (20 gradations). Therefore the inactivity of some

genes has a very small impact on survival, especially as each level of food availability occurs on average once in 20 generations. As a consequence, drift allows genes to become inactive. Furthermore, there is asymmetry in the effect of mutations. If a gene becomes inactive there may be no selection against the mutated rule for many generations and it will be weak if nearby genes are active. On the other hand, if a gene is switched on it takes the value of a nearby (lower γ) gene and so there is ini-

tially no selection for it, until it mutates to a lower value (since the optimal strategy is monotonically decreasing), whereas mutation to a higher value will be selected against. As a result, far fewer than all the genes remain switched on after many generations.

Of greater significance is the pattern of active genes that evolves under a small cost of flexibility ($\kappa = 0.05$, Fig. 2B). Deviations from optimality at low food availability (γ) have a bigger effect on survival than at high γ , which is why the evolved rule is more fine-grained at low γ (30% of maximum flexibility at $\gamma < 0.5$) than at high γ (6% of maximum flexibility at $\gamma \geq 0.5$). Hence we predict that under conditions where animals are attempting to use resources to survive, they will be more responsive to changes in the food supply when the food supply is poor than when it is rich. This occurs because of the shape of $x^*(\gamma)$, which decreases with γ at a decelerating rate. This in turn depends on the probability of going without food for a long time, which as a power function of $(1 - \gamma)$ is also a decelerating function. Thus, we expect this shape to apply generally to scenarios in which the main objective is to maximise survival, and so animals should be more responsive to changes in food availability when food is scarce. However, if the fitness function were not merely based on survival, but also for instance involved reproduction, its shape could be very different (McNamara et al., 1991), which might favour a different pattern of behavioural flexibility under the evolved rule. This should be the focus of future work.

The decision rule that evolves for $\kappa = 0.05$ uses only around four genes (20% of maximum flexibility) and so is reasonably crude, but despite this it suffers only a small reduction in survival: the probability of surviving is much closer to that achieved by the optimal strategy than it is under high flexibility costs. Thus, by adaptively reducing its flexibility to respond to changes in the foraging conditions, the evolved mechanism can achieve close to the same level of performance as the optimal strategy; any further improvements in fit to the optimal strategy would be outweighed by the costs of a more complex, flexible mechanism. We see that there is robust performance even when flexibility costs are high, which reduces the number of active genes but still allows the forager to perform almost as well as the optimum under most conditions. This results from the strong asymmetry in the cost of deviating from the optimal target level of reserves. If the animal attempts to store too much fat the reduction in survival is smaller than that from attempting to store too little fat. Note that

the values of active genes tend to lie *on* the optimal strategy curve (Fig. 2B), such that the threshold for inactive genes is typically slightly higher than optimal (note that the reverse is true when inactive genes take the value of the first active gene above them). In contrast, there is very little deviation from optimality in the direction of smaller reserves. This asymmetry is most clear when food is abundant (and so the optimal reserve level is low), and less prominent at low food availability (due to the stronger survival costs). Thus, the evolved strategy deviates from the optimum, but only in the direction that is least costly for survival. This is an example of errors in behavioural strategies being associated with low canonical cost (McNamara and Houston 1986; Houston et al., 1992)

Besides being less costly than the completely flexible optimal strategy, another key benefit of the evolved rule is that it generalises across similar conditions. The complex strategies found by dynamic programming cannot do this, because the values in each cell of the look-up table are, in principle, independent of each other; it is only possible to find the optimal decision in a given cell by calculating payoffs. By contrast, the evolved rule generalises across similar levels of food availability, and it does this in a sensible way: it can afford to generalise a lot at high food availability, and less so at low food availability. If strategies show very strong generalisation animals might not adjust their behaviour to changing food availability in a gradual, linear manner, but instead show marked changes at critical thresholds. In simple terms, they may do one thing if food is very scarce and another if food is abundant, with little sensitivity otherwise. Our work highlights that the extent to which decision rules generalise (as opposed to specialise) across conditions may itself be under selection (Ghirland and Enquist, 2007). In general, we argue that the ability to respond in a similar way to similar types of conditions, even those that have never been encountered previously, is likely to be a key property of many decision-making mechanisms in real organisms, and has implications for understanding animal preferences (Enquist and Johnstone, 1997)

However, there are likely to be some conditions in which a simple rule fails to perform well. One important consequence of inaccuracy in decisions for fat storage is obesity, which is a major health problem in many human societies (Prentice, 2001; Wells, 2006; Friedman, 2009). The level of food availability in the modern Western world is unprecedented during our evolutionary history (O'Dea, 1992). At very high levels of food avail-

lability, our model predicts that animals will store far more fat than is necessary. This is because storing too much fat is less costly than too little fat, and this asymmetry is strongest when food is abundant. We suggest that obesity may result from an approximate, suboptimal mechanism that has not evolved to control body weight accurately when food is super-abundant, leading to a level of fat storage that is much too high. This is consistent with the idea that obesity is common because feeding mechanisms do not function appropriately in the modern world (Prentice, 2001; Speakman, 2007). Furthermore, if the mechanisms controlling fat storage in humans and other animals have evolved to be fairly crude due to associated costs, sensitivity at low food availability is likely to be prioritised. Such a prioritisation is consistent with the common clinical observation that it is very difficult to lose weight on low-calorie diets (Heymsfield et al., 2007), to which any survival mechanism will be highly sensitive. Our work may therefore shed some light on both the incidence and persistence of obesity in human societies (Prentice, 2001; Wells, 2006; Friedman, 2009).

We have made an initial step towards studying the evolution of flexible mechanisms for behaviour. We have shown that incorporating the potential costs of flexibility may shed light on the errors made by real animals in making decisions. In doing so, we have illustrated a general principle about animal behaviour focussing on the paradigmatic ‘small bird in winter’ (Lima, 1986; McNamara and Houston, 1990; Houston and McNamara, 1993), but further developments should reflect other species and different scenarios. Elsewhere (Fawcett et al., 2014) we have highlighted the role that ecological complexity may play in influencing the strategies that animals use to make decisions. Genetic algorithms are a good way to investigate how mechanisms for behaviour might evolve: they simulate the constraints on what can evolve, such as the fitness of intermediate stages and the mutation-selection balance, whilst giving a solution that is not optimal but is likely to approximate optimality given these constraints (see Giske et al., 2013; Giske et al., 2014 for more complex examples). Genetic algorithms also allow greater flexibility in the types of mechanisms that can be explored. We hope that this work inspires further studies that incorporate mechanistic constraints into the theory of animal decision-making.

Acknowledgments The authors are grateful to John McNamara and Pete Trimmer for insightful discussion, and Jarl

Giske and two anonymous reviewers for comments on an earlier version of this manuscript. This work was supported by the European Research Council (Advanced Grant 250209 to A.I.H.) and a College for Life Sciences Fellowship at the Wissenschaftskolleg zu Berlin awarded to A.D.H.

References

- Brodin A, 2007. Theoretical models of adaptive energy management in small wintering birds. *Phil. Trans. Roy. Soc. B.* 362: 1857–1871.
- Clark CW, Mangel M, 2000. *Dynamic State Variable Models in Ecology*. Oxford, U.K.: Oxford University Press.
- Collins EJ, McNamara JM, Ramsey DM, 2006. Learning rules for optimal selection in a varying environment: Mate choice revisited. *Behav. Ecol.* 17: 799–809.
- DeWitt TJ, Sih A, Wilson DS, 1998. Costs and limits of phenotypic plasticity. *Trends Ecol. Evol.* 13: 77–81.
- Enquist M, Johnstone RA, 1997. Generalization and the evolution of symmetry preferences. *Proc. Roy. Soc. Lond. B.* 264: 1345–1348.
- Fawcett TW, Hamblin S, Giraldeau L-A, 2013. Exposing the behavioral gambit: The evolution of learning and decision rules. *Behav. Ecol.* 24: 2–11.
- Fawcett TW, Fallenstein B, Higginson AD, Houston AI, Mallpress DEW et al., 2014. The evolution of decision rules in complex environments. *Trends Cogn. Sci.* 18: 153–161.
- Friedman JM, 2009. Obesity: Causes and control of excess body fat. *Nature* 459: 340–342.
- Ghirlanda S, Enquist M, 2007. How training and testing histories affect generalization: A test of simple neural networks. *Phil. Trans. Roy. Soc. B.* 362: 449–454.
- Giske J, Eliassen S, Fiksen Ø, Jakobsen PJ, Aksnes DL et al., 2013. Effects of the emotion system on adaptive behavior. *Am. Nat.* 182: 689–703.
- Giske J, Eliassen S, Fiksen Ø, Jakobsen PJ, Aksnes DL et al., 2014. The emotion system promotes diversity and evolvability. *Proc. Biol. Sci.* 281: 20141096.
- Grimm V, Berger U, Bastiansen F, Eliassen S, Ginot V et al., 2006. A standard protocol for describing individual-based and agent-based models. *Ecological Modelling* 198: 115–126.
- Gosler AG, Greenwood JJD, Perrins C, 1995. Predation risk and the cost of being fat. *Nature* 377: 621–623.
- Hamblin S, 2013. On the practical usage of genetic algorithms in ecology and evolution. *Methods Ecol. Evol.* 4: 184–194.
- Heymsfield SB, Harp JB, Reitman ML, Beetsch JW, Schoeller DA et al., 2007. Why do obese patients not lose more weight when treated with low-calorie diets? A mechanistic perspective. *Am. J. Clin. Nutr.* 85: 346–354.
- Houston AI, McNamara JM, 1999. *Models of Adaptive Behaviour*. Cambridge, U.K.: Cambridge University Press.
- Houston AI, McNamara JM, Hutchinson JMC, 1993. General results concerning the trade-off between gaining energy and avoiding predation. *Philos. Trans. R. Soc. London Ser. B-Biological Sci.* 341: 375–397.
- Houston AI, McNamara JM, 1993. A theoretical investigation of the fat reserves and mortality levels of small birds in winter. *Ornis Scand.* 24: 205–219.

- Houston AI, McNamara JM, Thompson WA, 1992. On the need for a sensitive analysis of optimization models, or, this simulation is not as the former. *Oikos* 63: 513–517.
- Hutchinson JMC, Gigerenzer G, 2005. Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behav. Processes* 69: 97–124.
- Krams I, Cirule D, Suraka V, Krama T, Rantala MJ et al., 2010. Fattening strategies of wintering great tits support the optimal body mass hypothesis under conditions of extremely low ambient temperature. *Funct. Ecol.* 24: 172–177.
- Lee SJ, Witter MS, Cuthill IC, Goldsmith AR, 1996. Reduction in escape performance as a cost of reproduction in gravid starlings *Sturnus vulgaris*. *Proc. R. Soc. London Ser. B-Biological Sci.* 263: 619–623.
- Lima SL, 1986. Predation risk and unpredictable feeding conditions: Determinants of body-mass in birds. *Ecology* 67: 377–385.
- McNamara JM, 1990. The policy which maximizes long-term survival of an animal faced with the risks of starvation and predation. *Adv. Appl. Probab.* 22: 295–308.
- McNamara JM, Houston AI, 1986. The common currency for behavioral decisions. *Am. Nat.* 127: 358–378.
- McNamara JM, Houston AI. 1990. The value of fat reserves and the tradeoff between starvation and predation. *Acta Biotheor.* 38: 37–61.
- McNamara JM, Houston AI, 1997. Currencies for foraging based on energetic gain. *Am. Nat.* 150: 603–17.
- McNamara JM, Houston AI, 2009. Integrating function and mechanism. *Trends Ecol. Evol.* 24: 670–675.
- McNamara JM, Merad S, Houston AI, 1991. A model of risk-sensitive foraging for a reproducing animal. *Anim. Behav.* 41: 787–792.
- Meijer T, Mohring FJ, Trillmich F, 1994. Annual and daily variation in body mass and fat of starlings *Sturnus vulgaris*. *J. Avian Biol.* 25: 98–104.
- O’Dea K, 1992. Obesity and diabetes in “the land of milk and honey.” *Diabetes / Metab. Rev.* 8: 373–388.
- Prentice AM, 2001. Obesity and its potential mechanistic basis. *Br. Med. Bull.* 60: 51–67.
- Speakman JR, 2007. A nonadaptive scenario explaining the genetic predisposition to obesity: The “Predation release” hypothesis. *Cell Metab.* 6: 5–12.
- Wells JCK, 2006. The evolution of human fatness and susceptibility to obesity: An ethological approach. *Biol. Rev.* 81: 183–205.
- Witter MS, Cuthill IC, 1993. The ecological costs of avian fat storage. *Philos. Trans. R. Soc. London, Ser. B* 340: 73–92.
- Witter MS, Cuthill IC, Bonser RHC, 1994. Experimental investigation of mass-dependent predation risk in European starlings *Sturnus vulgaris*. *Anim. Behav.* 48: 201–222.