LUND UNIVERSITY

Bayes' theorem and its applications in animal behaviour

McNamara, J M; Green, R F; Olsson, Ola

Link to publication

Total number of authors:
3

Bayes' theorem and its applications in animal behaviour

John M. McNamara, Richard F. Green and Ola Olsson

Bayesian decision theory can be used to model animal behaviour. In this paper we give an overview of the theoretical concepts in such models. We also review the biological contexts in which Bayesian models have been applied, and outline some directions where future studies would be useful. Bayesian decision theory, when applied to animal behaviour, is based on the assumption that the individual has some sort of "prior opinion" of the possible states of the world. This may, for example, be a previously experienced distribution of qualities of food patches, or qualities of potential mates. The animal is then assumed to be able use sampling information to arrive at a "posterior opinion", concerning e.g. the quality of a given food patch, or the average qualities of mates in a year. A correctly formulated Bayesian model predicts how animals may combine previous experience with sampling information to make optimal decisions. We argue that the assumption that animals may have "prior opinions" is reasonable. Their priors may come from one or both of two sources: either from their own individual experience, gained while sampling the environment, or from an adaptation to the environment experienced by previous generations. This means that we should often expect to see "Bayesian-like" decision-making in nature.

J. M. McNamara, Dept of Mathematics, Univ. of Bristol, BS8 1TW, (john.mcnamara@bristol.ac.uk). – R. F. Green, Dept of Mathematics and Statistics, Univ. of Minnesota-Duluth, Duluth, MN 55812, USA. – O. Olsson, Dept of Animal Ecology, Lund Univ., Ecology Building, SE-223 62 Lund, Sweden.

Bayes' theorem

Bayes' theorem, which is named for Thomas Bayes, a Presbyterian minister and mathematician, who lived from 1702 to 1761, provides a method of determining probabilities, or parameters of probability distributions, based on observations. Roughly speaking, Bayes' theorem gives a method of calculating conditional probabilities. A Bayesian begins with a prior probability that some aspect of the world holds, then makes observations that modify that probability to produce a posterior probability. A familiar example involves a medical test. Imagine a disease that affects one percent of a population. A medical test is not completely accurate; it is positive for ninety percent of the people who have the disease and negative for eighty percent of people who do not have the disease. A person is chosen at random and given the test. The test is positive. What is the probability that the person has the disease? Out of 1 000 people, ten will have the disease, and nine of these would have a positive test. However, 990 people will not have the disease, but 20%, or 198 would have a positive test. Thus, $9+198$ people would have a positive test, but only 9 of these would have the disease. Therefore, the conditional probability of having the disease, given a positive test, is $9/(9+198) = 9/207 \approx 0.0435$. Before the test the probability of having the disease is 1%. This is the prior probability. After a positive test the probability is 4.35%. This is the posterior probability given the observation of a positive test. A similar calculation shows that the posterior probability of having the disease given a negative test is 0.00126, i.e. 0.126%.

More formally, Bayes theorem is as follows. Suppose that there are n possible states of the world, labelled

S_1, S_2, ..., S_n. The prior probability that S_i is the true state is $P(S_i)$. Let A be some event which has probability $P(A|S_i)$ of occurring given that S_i is the true state of the world. Then the overall (prior) probability that the event A occurs is

$$P(A) = P(A|S_1)P(S_1) + P(A|S_2)P(S_2) + \ldots P(A|S_n)P(S_n)$$

Given that the event A has been observed to have occurred the posterior probability that S_i is the true state of the world is

$$P(S_i|A) = \frac{P(A|S_i)P(S_i)}{P(A)}$$

Statistical inference is concerned with making inferences about unknown parameters from the observed results of some experiment whose outcome is (at least partly) random. For example it might be required to estimate the mean crop yield under a new fertiliser treatment from observations of the actual crop yields in a trial. One approach to statistical inference is known as Bayesian statistical inference. Bayes theorem is central to this approach. In the approach, before any observations are taken the possible values of an unknown parameter are given prior probabilities. Observations are then taken and these probabilities are subsequently modified via Bayes theorem to form the posterior probabilities. Inferences are then drawn from the posterior probabilities.

This paper, and others in this volume, are concerned with how animals integrate prior information and observations, so that Bayes theorem is also central to the theoretical considerations here. We are not, however, interested in what inferences can be made about unknown parameters, but rather whether animals make decisions that are optimal given the appropriate posterior probabilities. We are thus not concerned with Bayesian statistical inference, but with the related topic of Bayesian decision theory. For a discussion of how this framework can be applied to animal behaviour see McNamara and Houston (1980).

Statisticians differ over the merits of Bayesian and classical statistics. The main criticism of the Bayesian approach by classical statisticians concerns the use of prior probabilities by Bayesians. Classical statisticians ask Bayesians where they get their priors. The same criticism can be made of Bayesian decision theory. Although this may be a problem in some circumstances, we argue that it is not a problem when applying Bayesian decision theory to animal behaviour. As we detail below, it seems reasonable to assume that evolutionary history and previous experience determine well-defined priors, and then very reasonable that natural selection could produce animals that behave as if they knew these prior probabilities.

Where do priors come from?

There are two different processes that could lead an organism to behave as if it knew the prior distribution appropriate to its current environment.

Adaptation

If the ancestors of an organism have evolved in an environment in which the types of local habitat and their frequency of occurrence have been stable, then natural selection could lead to the organism behaving as if it knew this information. In fact to say that behaviour is adapted to the environment of its ancestors is essentially saying that the organism is using this prior information.

We can regard the 'worldview' of an organism as set by the environment experienced by its ancestors. This worldview may restrict what the animal is capable of learning. For example, animals may never learn that an environment is predator free, or never learn that they will not be interrupted while foraging. In Bayesian terms, the prior probability for these possibilities is zero, and since the prior is zero so is the posterior probability, no matter how strong the evidence to the contrary.

A restricted worldview could restrict the flexibility with which the organism responds to local conditions. For example, imagine a squirrel species whose members find themselves in one of two habitat types. In one habitat owls are the squirrels' sole predator, in the other snakes are the sole predator. If the ancestors of an individual only experienced owl predation, then the organism will have a rule about what are possible dangers, and how to deal with these dangers, that has been shaped by the danger of owls. This rule might be inappropriate when dealing with snakes, and the squirrel may never be able to learn the correct anti-predator behaviour against snakes. We can only expect current behaviour to be adapted to both sorts of predators if ancestors experience both in the past, some ancestors experiencing snakes and some owls. This illustrates a general point about phenotypic plasticity made by Houston and McNamara (1992).

The lack of flexibility may have profound consequences for experimental design. If an experiment involves a treatment that is entirely inconceivable to the animal, the response may not be the one predicted by the experimenter. The animal will be using a rule that is adapted to its natural environment but cannot respond in an adaptive way to the specific novel situation it now faces. Under these circumstances it is not possible to predict the response of the animal by considering the current situation alone; predictions must be based on prediction of the rule it should use in its natural environment (McNamara 1996). However, since many rules may do well in its natural environment, it may be

virtually impossible to predict behaviour in the novel situation from theoretical considerations.

Experience

An animal's prior distribution may also be determined by its past experience. For example, the animal might take the prior probability that it will rain on a particular day to be the frequency with which previous days have had rain.

In practice we can expect that the prior for any animal is typically set by a combination of the environment in which it evolved and its own past experience.

Biological examples

Example 1: foraging in patches

An animal is foraging in an environment in which food occurs in well-defined patches. Within each patch food occurs as discrete items, which the animal finds by searching the patch. Patches vary in the number of prey items they contain and maybe other characteristics such as the ease with which individual items are found. As the animal searches a patch it gains information about the characteristics of this patch from the number of items found so far and the times at which each item was found.

Suppose the animal has already had lots of experience in this environment during which it has gained information, where this information is equivalent to learning the types of patches present and their frequency of occurrence. The animal arrives at a new patch. Initially the probability that the patch is of any given type is equal to the frequency of that patch type in the environment as a whole. This is the prior probability. As the animal searches the patch it updates its estimate of patch type based on its experience. Examples that have been analysed include the following:

a) each patch contains either zero or one item. Then the longer the animal continues without finding an item, the more likely it is that the patch is empty (McNamara and Houston 1980, 1985a).

b) Patches differ in the number of prey items present. Each item in a patch is found after an exponential time, independently of the time to find other items. The mean of this exponential search time for each item is the same for all items and all patches. In this example, the current posterior probability that a patch contains a given number of items only depends on the prior distribution, the total number of items found on the patch so far and the total time to find these items (Oaten 1977, Green 1980, 1984, 1987, McNamara and Houston 1980,

McNamara 1982, Valone and Brown 1989, Olsson and Holmgren 1998).

c) Each patch contains exactly one item, but different patches vary in how hard the item is to find. An analogous problem is that of a squirrel cracking a nut – here the habitat is composed of different nuts, which differ in terms of how hard they are to crack. The squirrel tries to crack nuts sequentially. Thus each nut acts as a patch. The longer the time spent so far in an unsuccessful attempt, the more likely that this is a hard nut (McNamara and Houston 1985b, Green and Nuñez 1986).

We would expect the decisions of an animal while on the patch to depend on both experience on the current patch and previous experience in the environment. For example, consider an animal that maximises the rate at which it obtains food items in the environment. Then its decision to leave a patch should be influenced by the mean rate at which it can get items in other patches – the higher this rate the sooner it should leave the current patch. The decision to leave should also depend on future prospects for food on the current patch, which depends on posterior information about this patch. This information is determined by the animal's experience in the environment as a whole, which sets the prior, and experience on the patch (Green 1980, McNamara 1982, Olsson and Holmgren 1998).

Example 2: mate choice during an annual breeding season

Collins et al. (unpubl.) consider the following Bayesian model of mate choice. Suppose that each year each female member of a population must choose a male to mate with. Males vary in quality, and the female can determine the quality of a male by inspection. The female inspects a sequence of males, attempting to choose one of the highest quality males in the population as a mate. However, the distribution of male qualities varies from year to year, so that at the beginning of a breeding season, before a female has inspected any male, she does not know what range of qualities are high for that year. She does, however, have prior information on how the distribution of quality varies from year to year. For a semelparous species this information comes from the environment in which ancestors evolved. For an iteroparous species the female has this information and her experience in previous years. As the female inspects males during the current breeding season she updates her estimate of the distribution of quality this year.

To illustrate the updating process, suppose that within a given year male quality has a normal distribution with mean m and variance σ^2. Here the mean m varies from

year to year, although the within year variance σ^2 is the same each year. We suppose that the between year variation in the annual mean m has a normal distribution with mean μ_0 and variance v_0^2. For this scenario, as the female inspects males during a breeding season she gains information on the value of m for that year. At the start of the breeding season the female has only information that has been determined by her evolutionary history and her experience in previous years, so that the prior distribution of m is normal with mean μ_0 and the variance v_0^2. Suppose that later on in that breeding season the female has inspected a total of n males and found that the average quality of these males is $\bar{x}$. Then it can be shown using Bayes theorem that the posterior distribution of m given this information is normal with mean

$$\mu_n = (1 - \alpha_n)\mu_0 + \alpha_n \bar{x} \tag{1}$$

where

$$\alpha_n^{-1} = 1 + \sigma^2 / n v_0^2 \tag{2}$$

and variance

$$v_n^2 = \frac{v_0^2}{1 + n v_0^2 / \sigma^2} \tag{3}$$

(DeGroot 1970). As can be seen, the posterior mean μ_n is a weighted average of the prior mean μ_0 and the observed mean $\bar{x}$, with greater weight α_n given to the observed mean as the number of observations increases. The posterior mean μ_n provides an estimate of the true value of the mean m for that year. Not surprisingly, as the number of observations increases, these estimates tend to get better. This improvement can be seen from Eq. 3, which shows that the posterior variance of m decreases as n increases.

For further examples of a Bayesian approach to mate choice, see for example Luttbeg (1996, 2002) and Mazalov et al. (1996).

Example 3: growth under predation risk

Different individuals of a species are born into different types of environments. Environmental types differ in their predation risk. Before an individual has gained any information on the type of environment it is in, the prior probability that it is a particular type is the frequency with which its ancestors experienced this type. These prior probabilities are then updated to posterior probabilities in the light of the individual's experience in the environment. Observations that might provide useful information in this updating process include chemical cues as to the presence of predators (reviewed by Kats and Dill 1998), or the frequency with which predators are observed. Even without these obvious cues, the fact the organism is still alive gives it information – the

longer it has lived the lower the estimate of danger (Welton et al. 2003).

Mathematical calculation versus rules used by animals

Mathematical calculations using Bayes theorem are based on a characterisation of local environments into types, each of which has a specified prior probability of occurrence. We might expect natural selection to produce organisms that behave as if they know this information, but this does not mean that organisms characterise the environment in this way (McNamara and Houston 1980).

To illustrate the above point, consider the patch use example in which a nut forms a patch for a foraging squirrel. Suppose that the environment is composed of nuts of two distinct types; nuts that are easy to crack and nuts that are hard to crack. Three quarters of all nuts are easy. These nuts take an exponential time to crack with mean one minute. The remaining one quarter are hard, taking an exponential time to crack with mean five minutes. Under these assumptions the probability that a randomly selected nut takes more than time t to crack is

$$f(t) = \frac{3}{4} e^{-1} + \frac{1}{4} e^{-0.2t} \tag{4}$$

This is what the animal experiences, so this is what it is reasonable to assume it might learn. An animal may never learn there are two sorts of nuts, no matter how many it cracks. Nor does it need to know this information in order to be able to maximise the rate at which it cracks nuts. The distribution of time to crack a randomly selected nut has the property that the longer the time that has elapsed, the greater the further time it is likely to take. Specifically it can be shown that if the animal has tried to crack a nut without success for time t, then the expected further time taken is

$$1 + \frac{4}{3 e^{-0.8t} + 1} \tag{5}$$

This information is all that it is important for the animal to learn if it is to behave optimally. It need never know there are two distinct nut types, and consequently never characterise information in terms of prior and posterior probabilities.

Similar reasoning applies in the example with many patch types, each varying in the number of items per patch. In this case animals can be expected to learn how the number of items found so far and the current search time determine the likelihood of finding further items on the patch in the future. This information is all that is important if an animal is to forage optimally in this environment. The animal does not need to know there

are distinct patch types and their frequency of occurrence (McNamara and Houston 1987).

The above is concerned with implementation of the optimal rule. We would not, however, expect natural selection to produce organisms that are exactly optimal. Instead, we would expect them to use rules of thumb that do very well in the environment in which the species evolved. Often simple rules can perform surprisingly well and are highly robust (Houston et al. 1982, Gigerenzer and Todd 1999). These rules often require the use of much less information to implement than is required by the optimal rule. Thus in testing whether animals are 'Bayesian' we should be cautious. They are unlikely to be perfect Bayesians and their rules may only have Bayesian characteristics if there is significant selection pressure to adopt these characteristics.

Although we might expect an evolved rule to perform well in the environment to which an animal is adapted, the rule may perform badly if exposed to an unusual situation, for example in the laboratory (McNamara and Houston 1980, Houston and McNamara 1989, 1999).

The value of different sources of information

In some cases it may not matter that an animal has poor information about its local environment. To illustrate this, consider the mate choice example. Suppose that the mean quality of males in a year has high between year variance (i.e. high v_0^2.) This does not matter if the within year variance is small (i.e. small σ^2); in this case the female should not waste time inspecting males, but should choose the first male encountered because no other male is likely to be considerably better (Collins et al. in publ.).

In some cases only certain aspects of the information may be important. Consider the example of an animal foraging on patchy food. Suppose that the number of food items per patch is highly variable, but is always large. In this case the intake rate of an animal foraging on a patch might resemble a smooth flow that decreases over time. If this is so, the situation is similar to that envisaged by Charnov (1976), and an animal can approximately maximise its intake rate by leaving each patch when the flow rate falls to the average rate for the environment. In these circumstances it is likely to be important that an animal learns the average rate for the environment, but information about the frequency of different patch types is unimportant.

Suppose, however, that information on a particular aspect of the environment is valuable to an animal. How much weight should an organism put on observations of the current environment as opposed to prior information? Not surprisingly, the optimal weight depends on how specific is the prior information and how good the current observations. Equation 1–3 illustrate this for the mate choice example. The parameter α_n a gives the weighting of current observations. As can be seen from Eq. 2 the smaller the ratio of observation error to prior error (i.e. the smaller is σ^2/v_0^2) the bigger the weighting of current information. Equation 2 also shows that this weight increases as the number of observations n increases.

Some historical notes on foraging theory

Optimal foraging theory began with two papers published back to back in the American Naturalist in 1966 by MacArthur and Pianka (1966) and by Emlen (1966). These papers were concerned with the optimal choice of patches and prey items in patchy environments. One goal was to understand the population consequences of foraging behaviour. For example, if prey are scarce, predators should take a wider range of prey. Then, if competition is measured by diet overlap, competition should be more severe when prey are less abundant. Subsequent work in optimal foraging theory largely built on the concepts in these two papers.

One of the best known concepts in optimal foraging theory is the marginal value theorem (Charnov 1976). Charnov assumed that the environment is composed of well-defined food patches, where on each patch energy is obtained continuously at a rate that decreases with time in a patch. Charnov sought the strategy that maximised an animal's long-term rate of energy intake. He showed that under this strategy a forager should leave each patch when the rate of obtaining energy in that patch falls to highest possible long-term average rate of energy intake, γ. This marginal value theorem makes two qualitative predictions: (1) all patches should be left when the instantaneous rate of energy intake in them is the same, and (2) if travel time between patches is longer, perhaps because there are fewer patches in the environment, foragers should spend longer in each type of patch.

A central assumption of Charnov's model was that energy is obtained as a smooth flow, rather than as discrete food items at unpredictable times. Bayesian foraging arose during the heyday of optimal foraging theory, starting with a paper by Oaten (1977). This paper included stochastic foraging with discrete items and was written in response to Charnov's (1976) claim that his deterministic model for the patch residence–time problem could easily be made stochastic. Oaten tried to show that the introduction of patch variability and information use into foraging models is not trivial, and he showed that in some cases a forager that uses the information gained while searching patches may do much better than a forager that does not use information.

In Charnov's model a simple rule tells the animal when to leave a patch – it should leave when the

instantaneous rate falls to γ. When foraging is stochastic a version of this rule can be used provided that the animal always has complete information about patch quality. In particular, suppose that food items are equally hard to find, the animal knows the number of prey when arriving, how many are left at any given time, and searches at random. Then the expected instantaneous energy intake rate is known from the prey density. Hence, the need for a smooth gain function is obviated. The marginal value theorem will then specify when to leave a patch – when the instantaneous intake rate falls to γ. For discrete prey items, under suitable assumptions this amounts to leaving when the number of prey left is some fixed number (the giving up density), corresponding to that instantaneous intake rate (Brown 1988, Olsson and Holmgren 1999).

However, if the forager cannot be assumed to "know" the exact number of prey in the patch, it cannot use this omniscient rule. Krebs et al. (1974) suggested that animals might use a fixed giving-up-time rule to satisfy the marginal value theorem. That is, animals should leave a patch when they have gone some particular, fixed time (the "giving-up time") without finding a prey item in the patch. The idea of a fixed giving-up time is appealing, but it is not the optimal rule on which to base patch departure (Iwasa et al. 1981, Green 1984, 1987). Simply put, it is not optimal because it ignores all the information contained in the number of prey found, and the time spent searching for those prey, i.e. all the previous experience in the patch. It only uses the information contained in the time since the last capture.

Oaten (1977) assumed that a forager "knows" how prey are distributed in patches and what the travel time is between patches. The forager was assumed to use the patchleaving rule that maximized the long-term average rate of finding prey based on knowledge of the environment (prey distribution and travel time) and experience in a patch. Oaten described how to find the optimal rule in general, but the calculations were very complicated. Oaten worked out a couple of simple examples. More complicated examples may be solved by dynamic programming (Green 1980, 1987, Olsson and Holmgren 1998), but even these examples require some simplifying assumptions, in particular, that the proportion of a patch that has been searched by time t is a deterministic function of t.

Among the early papers in Bayesian foraging (Oaten 1977, Green 1980, Iwasa et al. 1981, McNamara 1982) that by Iwasa et al. (1981) is the one most cited, perhaps because the model presented is tractable and gives a very appealing rule for patch departure. The rule suggested is to leave all patches when the expected instantaneous intake rate falls to some constant level – that is applying the marginal value theorem to this expected rate rather than the actual rate when there is a smooth flow. What Iwasa et al. (1981) showed elegantly was that the

expected instantaneous intake rate could be found using only information on the searching time spent, and the number of prey found in the patch so far. That is, by using simple and available sampling information. Thus, this was a model that relaxed the assumption of complete information, but still worked in a stochastic setting. The problem with the model is that the rule suggested is only reasonable as long as finding an item does not improve the estimate of what is left in the patch (McNamara 1982). When the prey distribution is clumped (variance greater than mean), finding an item does improve this estimate. The rule is then no longer optimal (Green 1984, Olsson and Holmgren 1998). Olsson and Brown dwell on this clumped prey case at some length in this volume.

Evidence that animals are Bayesian

How can we identify whether animals are Bayesians? As so often when studying behaviour, we may have no means of establishing what is going on in the animal's mind per se. Thus, we will probably never be able to observe directly whether animals really do have mental constructs that represent prior and posterior probability distributions. In any case, we have argued that it is not necessary for an animal to have these constructs in order to exhibit Bayesian type behaviour. Such behaviour can result from the animal using simple rules. In testing whether animals are Bayesian we are therefore not concerned with an animal's mental constructs, but with comparing the behaviour of an animal with that predicted by a model assuming Bayesian information use, and that predicted by alternative models.

The central feature of a Bayesian model is the dependence of present behaviour on prior information and current experience. So Bayesian behaviour could possibly be inferred if there is evidence that both of these components influence the behaviour of organisms in an adaptive way. Strong evidence for such Bayesian behaviour would have to show that if either is held fixed and the other altered, then this alters behaviour appropriately. There are, however, some problems with this simple criterion.

If the prior is set by experience then evidence for the effect of a prior can be sought by experimentally altering this prior experience. If, however, the prior is set partly by evolutionary history, then this aspect of the prior cannot be altered by experimental procedures. Instead, indirect evidence of the effect of the prior has to be obtained by comparative studies that look at populations or species with different evolutionary histories.

Just because there is evidence that animals are affected by current experience it does not necessarily mean that they are learning. Collins et al. (unpubl.) give a hypothetical example from mate choice. As in example

2, suppose that the distribution of male qualities varies from year to year. In a particular year a female inspects males sequentially and hence gains information on this year's distribution. She must choose a mate without recall of previously rejected males. Under the optimal Bayes rule the decision to accept a male takes into account the qualities of all males so far observed. But the female can do very well by just employing a decreasing acceptance threshold that depends on time alone. This rule has many of the properties one would associate with a learning rule – if the males observed so far are poor quality the female does not accept any but continues searching with a reduced threshold. She is thus less choosy, and so behaves as if she has learnt that this year is a poor year and responds appropriately. But to what extent can we describe a female that employs this simple deterministic threshold rule as learning?

Evidence from foraging

So what is the evidence that animals exploiting patchily distributed food are Bayesian foragers? In this case an animal learns about the environment from its previous experience. One immediate question is what sort of information does it learn? Does it, for example, just learn some simple measure of overall environment quality, or does it learn more subtle aspects of the environment.

In some cases only a simple overall measure of environmental quality is relevant. For example, in the patch use scenario envisaged by Charnov (1976) the environment is composed of food patches that deliver food rewards as a smooth decreasing flow. The simplest measure of overall environmental quality in this setting is the maximum long-term reward rate γ. We might regard learning this quantity as a Bayesian problem in itself, with the prior set from evolutionary history or previous environments encountered. Certainly there are simple Bayes-like rules that an animal can employ to learn γ given sufficient time (McNamara 1985, McNamara and Houston 1985a). Having learnt γ the animal maximises its rate of energy gain by leaving each patch when the local rate falls to γ. Here the animal is responding to both its previous experience in the environment and the current patch. But this is only evidence of Bayesian behaviour in some weak sense, since the animal only needs to know one aspect of the environment (γ) and knows this prior aspect with certainty when encountering a new patch. Some of the papers in this volume only provide evidence for Bayesian foraging in this sense.

In other cases more detailed knowledge of the environment is important if the animal is to behave optimally on a patch. In particular it is crucial whether the distribution of prey items is clumped with some

patches containing few items and others many. There is certainly evidence that animals do make adaptive use of this information and other pertinent information in a number of cases (Valone and Brown 1989, van Gils et al. 2003) as is reviewed by Valone (Valone 2006) in this volume.

There are some additional difficulties in generating testable predictions, which have still not been fully addressed. First, it may sometimes be difficult to identify the alternative model in a relevant manner. Of course, the choice of alternative model influences the choice of predictions to test. Second, the Bayesian foraging models to date assume that animals use no other source of information besides the prior and sampling information. However, that may not necessarily be the case. These two issues are partly interwoven.

Other sources of information could influence behaviour in a Bayesian manner, or interact with Bayesian behaviour. For example sensory information could act in this way. Patch foraging models usually assume that the animal either uses only sensory information (Valone and Brown 1989), or only sampling information (Green 1980, McNamara 1982) to guide patch departure. Some animals, such as rodents, may have olfaction that is good enough to determine patch quality quite accurately (Valone and Brown 1989). However, a combination of sensory information and sampling information is plausible. The sensory information may then act to modify the prior, and this could be a continuous process throughout the patch visit. So far, this possibility has not been exhaustively considered in any theoretical or empirical work. If sensory information is used to modify the prior at the arrival in the patch, it can be seen as classifying patches into types (sensu Stephens and Charnov 1982, Stephens and Krebs 1986). Sampling can then determine the subtype, i.e. qualities of the patch that cannot be determined beforehand.

A forager able to sense patch quality entirely accurately is able to follow the discrete version of the marginal value theorem and leave all patches at the same giving-up density (Brown 1988). A completely ignorant forager, incapable of using any information of patch quality, should leave all patches after a fixed searching time. However, simply observing some deviation from either or both of these predictions is not really evidence of Bayesian foraging. Specific predictions pertaining to the given system should be tested, such as higher giving-up densities (GUDs) in rich patches than in poor (Valone and Brown 1989), a positive relation between GUDs and searching time (Olsson and Holmgren 1999, Olsson et al. 1999), or a combination of a rejection of alternative models and a lack of relation between estimated potential intake rate at departure and initial prey density (van Gils et al. 2003).

Mate choice

Here there is certainly evidence that females searching for suitable mates adjust their acceptance thresholds depending on the phenotypes of previously encountered males (Jennions and Petrie 1997). This is certainly consistent with Bayesian mating behaviour.

Partial reinforcement extinction effect

Consider a Skinner box experiment in which an animal can obtain rewards by pressing a lever. Not all responses are rewarded. The experiment is set up so that, up to a certain time all responses are rewarded with probability p, after this time no responses are rewarded. The animal has no information on this extinction time other than through the rewards it receives. The empirical finding from this experiment is that, after extinction, animals stop responding more quickly when p is high than when it is low (reviewed by Mackintosh 1974).

McNamara and Houston (1980) put a Bayesian interpretation on this finding. In the natural environment food sources always run out eventually. The animal has a prior probability distribution on the length of time a new source will last. It combines this with the reward information on the source to find the posterior probability that the source has extinguished. The larger is p, the more likely that a run of unrewarded responses is due to extinction rather than a run of bad luck, and the greater the posterior probability that extinction has occurred.

Future directions

We think it is worth widening the areas to which Bayes theorem has been applied. Obvious areas outside foraging are further applications to mate choice and predation risk.

Within, as well as outside of the realm of foraging studies, we think there has not been sufficient attention paid to the effect of the evolutionary environment on prior distributions. We give three simple examples.

(i) Most models of foraging do not differentiate between animals on the basis of their evolutionary environment, treating all animals as capable of learning prior distributions from experience. But, for instance, it may be that only animals whose ancestors have encountered clumped prey distributions are capable of learning when prey are clumped and responding accordingly.

(ii) An animal whose ancestors never lived in a predator-free environment may not be able to learn that its current environment (e.g. in a lab) is safe; whereas others with a different evolutionary history may be capable of learning this.

(iii) One reason for preferring immediate to delayed rewards is that an animal may be interrupted and lose the reward if it does not take what is available now. Animals of different species may have different preferences for immediacy as a consequence of their ancestors experiencing different interruption rates.

In addition to gaps in knowledge identified above, it appears that more work is needed regarding how groups of animals form decisions in a Bayesian context (Valone and Giraldeau 1993, Sernland et al. 2003). Also in the field of habitat selection there is a need to evaluate how Bayesian information processing alters predictions. A final remaining issue (MacArthur and Pianka 1966) that has only rarely been touched upon (Green 1990, Rodríguez-Gironés and Vásquez 1997, Olsson and Holmgren 2000) is how Bayesian behaviour influences population and community dynamics.

References

Brown, J. S. 1988. Patch use as an indicator of habitat preference, predation risk, and competition. – Behav. Ecol. Sociobiol. 22: 37–47.

Charnov, E. L. 1976. Optimal foraging, the marginal value theorem. – Theor. Popul. Biol. 9: 129–136.

DeGroot, M. H. 1970. Optimal statistical decisions. – McGraw-Hill.

Emlen, J. M. 1966. The role of time and energy in food preference. – Am. Nat. 100: 611–617.

Gigerenzer, G. and Todd, P. M. 1999. Simple heuristics that make us smart. – Oxford Univ. Press.

Green, R. F. 1980. Bayesian birds: a simple example of Oaten's stochastic model of optimal foraging. – Theor. Popul. Biol. 18: 244–256.

Green, R. F. 1984. Stopping rules for optimal foragers. – Am. Nat. 123: 30–40.

Green, R. F. 1987. Stochastic models of optimal foraging. – In: Kamil, A. C., Krebs, J. R. and Pulliam, H. R. (eds), Foraging behavior. Plenum Press, pp. 273–302.

Green, R. F. 1990. Putting ecology back into optimal foraging theory. – Comm. Theor. Biol. 1: 387–410.

Green, R. F. and Nuñez, A. T. 1986. Central place foraging in a patchy environment. – J. Theor. Biol. 123: 35–43.

Houston, A. I. and McNamara, J. M. 1989. The value of food: effects of open and closed economies. – Anim. Behav. 37: 546–562.

Houston, A. I. and McNamara, J. M. 1992. Phenotypic plasticity as a state-dependent life-history decision. – Evol. Ecol. 6: 243–253.

Houston, A. I. and McNamara, J. M. 1999. Models of adaptive behaviour. – Cambridge Univ. Press.

Houston, A. I., Kacelnik, A. and McNamara, J. M. 1982. Some learning rules for acquiring information. – In: McFarland, D. J. (ed.), Functional ontogeny. Pitman, pp. 140–191.

Iwasa, Y., Higashi, M. and Yamamura, N. 1981. Prey distribution as a factor determining the choice of optimal foraging strategy. – Am. Nat. 117: 710–723.

Jennions, M. D. and Petrie, M. 1997. Variation in mate choice and mating preferences: a review of causes and consequences. – Biol. Rev. 72: 283–327.

Kats, L. B. and Dill, L. M. 1998. The scent of death: chemosensory assessment of predation risk by prey animals. – Ecoscience 5: 361–394.

Krebs, J. R., Ryan, J. C. and Charnov, E. L. 1974. Hunting by expectation or optimal foraging-study of patch use by chickadees. – Anim. Behav. 22: 953–964.

Luttbeg, B. 1996. A comparative Bayes tactic for mate assessment and choice. – Behav. Ecol. 7: 451–460.

Luttbeg, B. 2002. Assessing the robustness and optimality of alternative decision rules with varying assumptions. – Anim. Behav. 63: 805–814.

MacArthur, R. H. and Pianka, E. R. 1966. On optimal use of a patchy environment. – Am. Nat. 100: 603–609.

Mackintosh, N. J. 1974. The psychology of animal learning. – Academic Press.

Mazalov, V., Perrin, N. and Dombrovsky, Y. 1996. Adaptive search and information updating in sequential mate choice. – Am. Nat. 148: 123.

McNamara, J. M. 1982. Optimal patch use in a stochastic environment. – Theor. Popul. Biol. 21: 269–288.

McNamara, J. M. 1985. An optimal sequential policy for controlling a Markov renewal process. – J. Appl. Prob. 22: 324–335.

McNamara, J. M. 1996. Risk-prone behaviour under rules which have evolved in a changing environment. – Am. Soc. Zool. 36: 484–495.

McNamara, J. M. and Houston, A. I. 1980. The application of statistical decision theory to animal behaviour. – J. Theor. Biol. 85: 673–690.

McNamara, J. M. and Houston, A. I. 1985a. A simple model of information used in the exploitation of patchily distributed food. – Anim. Behav. 33: 553–560.

McNamara, J. M. and Houston, A. I. 1985b. Optimal foraging and learning. – J. Theor. Biol. 117: 231–249.

McNamara, J. M. and Houston, A. I. 1987. Foraging in patches: there's more to life than the marginal value theorem. – In: Commons, M. L., Shettleworth, S. and Kacelnik, A. (eds), Quantitative analyses of behaviour. Vol. VI. Lawrence Erlhbaum, pp. 23–39.

Oaten, A. 1977. Optimal foraging in patches: a case for stochsticity. – Theor. Popul. Biol. 12: 263–285.

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

Olsson, O. and Holmgren, N. M. A. 1999. Gaining ecological information about Bayesian foragers through their behaviour. I. Models with predictions. – Oikos 87: 251–263.

Olsson, O. and Holmgren, N. M. A. 2000. Optimal Bayesian foraging policies and prey population dynamics – some comments on Rodríguez-Gironés and Vásquez. – Theor. Popul. Biol. 57: 369–375.

Olsson, O., Wiktander, U., Holmgren, N. M. A. et al. 1999. Gaining ecological information about Bayesian foragers through their behaviour. II. A field test with woodpeckers. – Oikos 87: 264–276.

Rodríguez-Gironés, M. A. and Vásquez, R. A. 1997. Density-dependent patch exploitation and acquisition of environmental information. – Theor. Popul. Biol. 52: 32–42.

Sernland, E., Olsson, O. and Holmgren, N. M. A. 2003. Does information sharing promote group foraging? – Proc. R. Soc. B 270: 1137–1141.

Stephens, D. W. and Charnov, E. L. 1982. Optimal foraging: some simple stochastic models. – Behav. Ecol. Sociobiol. 10: 251–263.

Stephens, D. W. and Krebs, J. R. 1986. Foraging theory. – Princeton Univ. Press.

Valone, T. J. 2006. Are animals capable of Bayesian updating? An empirical review. – Oikos 000: 000–000.

Valone, T. J. and Brown, J. S. 1989. Measuring patch assessment abilities of desert granivores. – Ecology 70: 1800–1810.

Valone, T. J. and Giraldeau, L.-A. 1993. Patch estimation by group foragers: what information is used? – Anim. Behav. 45: 721–728.

van Gils, J. A., Schenk, I. W., Bos, O et al. 2003. Incompletely informed shorebirds that face a digestive constraint maximize net energy gain when exploiting patches. – Am. Nat. 161: 777–793.

Welton, N. J., McNamara, J. M. and Houston, A. I. 2003. Assessing predation risk: optimal behaviour and rules of thumb. – Theor. Popul. Biol. 64: 417–430.