Behavioural Processes 89 (2012) 86–94

Contents lists available at SciVerse ScienceDirect

Behavioural Processes

journal homepage: www.elsevier.com/locate/behavproc

Tracking a changing environment: Optimal sampling, adaptive memory and overnight effects

Aimee S. Dunlap*, David W. Stephens

Ecology, Evolution & Behavior, University of Minnesota, St. Paul, MN, USA

ARTICLE INFO

Article history:
Received 3 May 2011
Received in revised form 4 October 2011
Accepted 8 October 2011

Keywords:
Tracking
Sampling
Memory
Environmental change
Foraging

ABSTRACT

Foraging in a variable environment presents a classic problem of decision making with incomplete information. Animals must track the changing environment, remember the best options and make choices accordingly. While several experimental studies have explored the idea that sampling behavior reflects the amount of environmental change, we take the next logical step in asking how change influences memory. We explore the hypothesis that memory length should be tied to the ecological relevance and the value of the information learned, and that environmental change is a key determinant of the value of memory. We use a dynamic programming model to confirm our predictions and then test memory length in a factorial experiment. In our experimental situation we manipulate rates of change in a simple foraging task for blue jays over a 36 h period. After jays experienced an experimentally determined change regime, we tested them at a range of retention intervals, from 1 to 72 h. Manipulated rates of change influenced learning and sampling rates: subjects sampled more and learned more quickly in the high change condition. Tests of retention revealed significant interactions between retention interval and the experienced rate of change. We observed a striking and surprising difference between the high and low change treatments at the 24 h retention interval. In agreement with earlier work we find that a circadian retention interval is special, but we find that the extent of this 'specialness' depends on the subject's prior experience of environmental change. Specifically, experienced rates of change seem to influence how subjects balance recent information against past experience in a way that interacts with the passage of time.

1. Introduction

Animals live in a changing world. Some change occurs predictably (e.g. daily and seasonal patterns) but much of this change occurs unpredictably. Animals often need to respond to this change. For example, they must adjust their behavior to a change in prey quality or the presence of a predator. Adjusting to a changing world is a two part problem. The animal must somehow use experience to obtain information about the current state of the world (which we call 'tracking'), and the animal must somehow retain this information and translate it into appropriate action. Taken together these two aspects of the 'adjusting to change' problem call upon nearly every aspect of an animal's cognitive machinery from perception to decision-making. As a crude caricature however, these two functional processes crudely coincide with learning and memory. When the animal tracks the environment, it learns about the states of the environment as it experiences changed conditions. Similarly, when animal later acts on this in response to this experience, it must call upon some 'stored' or remembered representation of this experience. In practice, we cannot separate these two aspects of the problem. We cannot investigate the retention of information unless the animal has acquired it previously; and we cannot investigate the acquisition of information unless the animal retains it in some way. The intimate relationship between learning and memory will be familiar to many students of animal psychology. Yet, as a matter of research strategy, it is often convenient to focus on either the acquisition or retention side of the problem in any given study.

For example, several models have been used to apply ideas from behavioral ecology and foraging theory to the acquisition side of this problem under the general heading of environmental tracking (reviewed in Stephens, 2007; Stephens and Dunlap, 2008). These models are typically used to ask how frequently an animal should 'sample' the resources in its environment to optimally adjust its behavior to changing conditions. Obviously enough, the degree and nature of environmental change is a fundamental part of these models. Qualitatively, both low and high rate of change can reduce the value of environmental tracking. If the environment never changes, the animal can adapt to it genetically and

---

* Corresponding author. Current address: Ecology & Evolutionary Biology, 1041 E Lowell St, University of Arizona, Tucson, AZ 85721, USA.

E-mail addresses: asdunlap@arizona.edu, asdunlap@email.arizona.edu (A.S. Dunlap).

0376-6357/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.beproc.2011.10.005

so it does not need a learning-like mechanism. On the other hand, if the environment changes unpredictably there is literally nothing to track, and again the animal does not need a tracking mechanism. Between these two extremes, we expect that higher rates of change will necessitate higher sampling effort. Theoretically, the relative costs of errors of sampling too much (checking resources too frequently in a stable environment) and sampling too little (checking too infrequently and missing the opportunity to exploit a profitable environmental change) interact with the rate of change in the environment to determine the optimal degree of sampling effort. This body of theory is relatively mature, and we now have several experimental studies of these models (see Stephens, 2007 for a review). Curiously, most of these experiments have manipulated the theoretically important 'cost of error' variables, and not the more fundamental environmental rate of change.

In contrast to the acquisition side of the problem, behavioral ecologists have paid relatively little attention to the output or recall side of the tracking problem. A few models have addressed this problem by modeling the fitness value of memory, typically linking this to environmental variability in a relatively obvious way: long-lasting memory does not pay when environments change quickly because information go 'out of date' when things change quickly (e.g. Anderson and Schooler, 1991; Dunlap et al., 2009). These models commonly consider how a memory system should integrate old and new experience. In doing this, they use a range of mathematical weighting schemes (e.g. exponential weighting, linear operators or Bayesian analysis) to explore past versus recent balance (e.g. Devenport et al., 1997; Harley, 1981; Houston et al., 1982; Killeen, 1981; Valone, 2006). These 'weighted memory' models commonly predict that animals should give more weight to recent experience in changing environments. The rationale for this result clearly follows the same logic as in memory length models: in changing environments old information becomes 'out-dated' quickly (Mangel, 1990; McNamara and Houston, 1985, 1987). Empirical work that addresses the fitness value of memory is relatively rare, although there are many snippets within the larger literature of learning and memory that bear on it in one way or another. A simple and elegant series of studies by Devenport and Devenport shows how empirical studies might proceed by testing how animals weight new versus old information, and qualitatively supports the ideas developed within this paper (Devenport and Devenport, 1994; Devenport et al., 1997).

Within these models, change is an abstraction. They can accommodate change occurring at arbitrary time scales. In nature, however, some time scales are more important than others. Daily and annual cycles, for example, surely impose some structure on the types of changes that matter to animals in nature. Many readers will know that periods of 24 h are special in several ways, both physiologically and behaviorally. Animals seem to recall events that happen at 24 h intervals more accurately in some experimental situations (e.g. Daan, 2000; Pahl et al., 2007; Prabhu and Cheng, 2008a; Zhang et al., 2006). Yet, models of tracking and adaptive memory do not incorporate these special intervals. In practice this may be because experimental studies of tracking focus on relatively short time intervals (i.e. less than a day), but surely animals face the problem of tracking at many different time scales so a general exploration of this problem must consider both long and short-term environmental changes. Consider the difference between an animal adapted to high rates of change and another adapted to low rates of change. Even without a formal model, we expect that animals in environments with high rates of change should track their environments relatively closely, and they should respond to changes relatively quickly. Animals in high-change conditions should also devalue old information, because high rates of change mean the information becomes outdated quickly. In contrast, animals adapted to lower change rates should adjust to changed condition more slowly and rely more heavily on old information. This paper tests these intuitive claims about the acquisition and retention of information via dynamic programming model and a straightforward factorial experiment. We first modeled our specific experimental scenario to confirm the general predictions from the literature, and then tested captive blue jays in two change regimes (high and low change rates). To assess how these rates of change influenced the jay's sensitivity to aging information, we tested five different retention intervals, creating a 2 by 5 factorial experiment with two levels of change rate, and five levels of retention interval.

Table 1
Possible outcomes concerning the state variable.

Event	Probability
A true and *A* recognized	$p(1-\varepsilon)$
A true and *B* recognized	$p\varepsilon$
B true and *A* recognized	$(1-p)\varepsilon$
B true and *B* recognized	$(1-p)(1-\varepsilon)$

2. Model

We first model a simple scenario considering an animal choosing between two stimuli, with this choice being repeated across a series of trials (time steps). At any given trial, the subject can select stimulus *A* or stimulus *B*. Only two situations are possible: *A* correct and *B* wrong, or *A* wrong, *B* correct. We call these the *A* correct and *B* correct states. We assume that a symmetric persistence process governs the transition between these two states, with q giving the probability of remaining in the same state (and $1-q$ the probability of switching correct states). In the absence of complete unpredictability, it would seem that there is nothing to track in this situation because an animal can immediately recognize which state is true. For instance, if the subject samples stimulus *A* and fails to obtain a reward, then it should immediately switch to choosing stimulus *B* (because only one stimulus can be correct at a time). It is a simple, but likely unrealistic, result that assumes animals can perfectly recognize the current state (*A* or *B* correct), and this is a difference from previous models. Instead, we distinguish between the true state and the recognized state. We do this using an error rate, ε, which we define as the probability that the recognized state differs from the true state.

We solved this problem using dynamic optimization, first determining the optimal behavior in the final time step, and then using this value to determine the optimal behavior for each prior time step. Let p be the subjective probability of the *A* correct state. In the final time step, then, the optimal behavior is to choose *A* if $p \geq ½$, with a payoff of p, and to choose *B* if $p < ½$, with a payoff of $1-p$. The optimal payoff for the last time step is a function of p, which we will call $V_0(p)$. This takes the form of a simple V shape.

Solving the second to last step is a bit more complicated. The payoffs are the same (choose *A*, get p; choose *B*, get $1-p$). Regarding the state variable, p, four things can happen: either *A* or *B* could be true, and each could either be correctly or incorrectly recognized (Table 1):

The animal, however can only subjectively observe two events: recognizing *A* or recognizing *B*, and these occur regardless of the actual choice the subject makes (Table 2).

If *A* is recognized, we can find the probability that *A* is true by Bayes Theorem:

Table 2
Events observable by the subject.

Event	Probability
A recognized	$p(1-\varepsilon)+(1-p)\varepsilon$
B recognized	$p\varepsilon+(1-p)(1-\varepsilon)$