Animal Behaviour 138 (2018) 59–67

Contents lists available at ScienceDirect

Animal Behaviour

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

The adaptive significance of age-dependent changes in the tendency of individuals to explore

Thomas N. Sherratt [a,*], Julie Morand-Ferron [b]

[a] Department of Biology, Carleton University, Ottawa, ON, Canada
[b] Department of Biology, University of Ottawa, Ottawa, ON, Canada

ARTICLE INFO

Article history:
Received 8 August 2017
Initial acceptance 15 September 2017
Final acceptance 4 December 2017

MS. number: A17-00639R2

Keywords:
ageing
bandit model
Bayesian learning
dynamic programming
exploration–exploitation
neophilia
neophobia
novelty

Many organisms show a reduced tendency to investigate unfamiliar objects as they age. Although the phenomenon could arise for a range of reasons, it is possible that this age-related increase in conservatism is adaptive. In particular, we propose that novel objects encountered late in life will be perceived as being relatively rare, so the value of information from investigating their properties will be estimated to be low. In addition, agents that investigate novel objects late in their lives will have little time left to exploit this information should the objects turn out to be profitable. We formalize the above arguments by developing an exploration–exploitation ('bandit') model. In this model, agents must decide whether to explore or ignore a novel object that it has just encountered at a given stage in its life, despite uncertainty regarding the commonness of the object in the environment and the likelihood that the object is profitable/unprofitable. We assume that, as agents encounter (and possibly investigate) unfamiliar objects, they use Bayesian inference to update their beliefs about the objects' commonness and profitability. Dynamic programming is concurrently used to identify the conditions under which the agent should explore or ignore these objects. If the benefit/cost ratio of investigating novel objects is high, then all individuals will be selected to explore regardless of their age. Likewise, if the ratio is low, then all individuals should ignore novel objects. Under intermediate conditions, young individuals that encounter novel objects should investigate their properties, while older ones should ignore them. The optimal switch in strategy arises as a consequence of age-dependent variation in both the novel object's perceived abundance and the future value of information regarding the object's profitability. We highlight several additional testable predictions of the model and discuss alternative adaptive explanations.

Many animals show an aversion to investigating the properties of objects they have never encountered before, a behaviour known as neophobia (Greggor, Thornton, & Clayton, 2015). Intriguingly, there is evidence that the tendency of individuals to exhibit such neophobia increases with age, while the tendency to spontaneously explore novel objects, termed neophilia, decreases with age. For example, immature wild baboons (*Papio ursinus*) and geladas (*Theropithecus gelada*) were found to exhibit more neophilia towards novel objects (a doll, ball or can) than conspecific adults (Bergman & Kitchen, 2009). Likewise, juvenile chimango caracara, *Milvago chimango*, show a greater tendency to investigate novel plastic objects and less neophobia than adults (Biondi, Bo, & Vassallo, 2010, 2013, 2015). When presented with baited puzzle-boxes, wild juvenile spotted hyaena, *Crocuta crocuta*, were more exploratory, less neophobic and more persistent than adults (Benson-Amram & Holekamp, 2012; Benson-Amram, Weldele, & Holekamp, 2013). Similarly, juvenile great tits, *Parus major*, and blue tits, *Cyanistes caeruleus*, visited novel problem-solving devices more frequently than adult conspecifics (Morand-Ferron, Cole, Rawles, & Quinn, 2011, 2015). Laboratory rodents also often demonstrate age-related declines in exploratory tendency from maturity onwards (e.g. longitudinal designs: Joyal et al., 2000; cross-sectional designs: Lamberty & Gower, 1990), while extensive surveys indicate that the openness of humans to new experiences tends to decline (albeit slightly) from middle age (e.g. Costa et al., 1986; Donnellan & Lucas, 2008; Roberts, Walton, & Viechtbauer, 2006).

Importantly, not all studies have revealed a decline in the tendency to explore with age (e.g. see Herborn et al., 2010; Hopper et al., 2014; Kendal, Coe, & Laland, 2005), and sometimes the relationship between age and exploratory tendency is complex. For

* Correspondence: T. N. Sherratt, Department of Biology, Carleton University, Ottawa, ON K1S 5B6, Canada.
E-mail address: Tom.Sherratt@Carleton.ca (T. N. Sherratt).

https://doi.org/10.1016/j.anbehav.2018.01.025
0003-3472/

example, 2-month-old mound-building mice, *Mus spicilegus*, showed a higher willingness to investigate a novel object (a white ball) compared to 12-month-old adults, but mice at 24 months of age behaved in a similar way to the juveniles (Lafaille & Feron, 2014). Nevertheless, despite these exceptions, there appears to be a general pattern of reduced exploratory tendencies in a wide range of species as individuals get older (Greenberg & Mettke-Hofmann, 2001; Mata, Wilke, & Czienskowski, 2013; O'Hara et al., 2017). Of course, the capacity of individuals to learn and solve problems is conceptually distinct from the willingness to approach novel stimuli (Griffin & Guez, 2014), but the adage 'you can't teach an old dog new tricks' (or indeed an old parrot, see Loepelt, Shaw, & Burns, 2016) is likely in part a reflection of the relative impassiveness and conservatism of adults (Menzel, 1966).

Why would individuals show an age-related decline in exploratory tendency? Whether the changes are adaptive or not, younger individuals may not have learned that the world is a risky place and may exhibit a degree of naiveté not seen in adults. One could also argue that this age dependency is merely a reflection of some other factor associated with age. In hierarchical societies for instance, younger individuals may have a subordinate role that makes them more exploratory because they do not have the access to resources that more dominant individuals do (Reader & Laland, 2001). Here we consider a simple adaptive explanation for the decline in exploratory tendency that does not invoke any of these mechanisms. Specifically, we argue that information about new things will be of less future value to individuals as they age. Therefore, younger individuals should be more prepared than older individuals to accept risks associated with investigating novel objects.

Intuitively, one might expect that organisms are selected to engage in behaviours that help them gain information early in life (exploration) and then, having learned much of what is to be learned, switch to using this information late in life (exploitation) (see Mata et al., 2013). However, while exploration–exploitation models have been widely invoked to investigate the trade-off between gaining information and using it (Cohen, McClure, & Yu, 2007; Mata & von Helversen, 2015; Reader, 2015), remarkably few models have been developed to understand age-related investment in information acquisition. The model of Eliassen, Jorgensen, Mangel, and Giske (2007) has long served to highlight the general influence of age (and life span) in the exploration–exploitation trade-off. Unfortunately however, their model has limitations when applied to elucidating the adaptive significance of any age-dependent tendency to investigate unfamiliar objects. This is because Eliassen et al. (2007) framed their model in terms of patch foraging and contrasted the performance of learners with nonlearners, rather than asking directly when an agent of given age should explore or ignore a newly encountered object. Berger-Tal, Nathan, Meron, and Saltz (2014) developed a more general exploration–exploitation model that identified four key stages of knowledge acquisition and use over the agent's life span, culminating in exploitation. While elegant, their model was not developed specifically to understand age-dependent changes in the tendency of individuals to explore novel objects.

In this paper we present a tailor-made model that directly identifies the optimal (long-term payoff maximizing) strategy of an agent that encounters an unfamiliar object at a given stage in its life. On first and subsequent encounter(s) with an object, the agent must simply decide whether to take the risk and explore the unfamiliar object's properties (for a potential immediate cost or benefit, as well as information it could use in the future), or ignore it altogether. Assuming that the agents learn about the abundance and profitability of the novel object in a Bayesian manner, we can identify the optimal decision of the agent by combining the agent's current information and expected payoff with the expected future value of information. Formalizing the phenomenon in this way not only allows us to elucidate the agent's optimal exploration–exploitation strategy, but also to isolate the role of two correlated properties of an agent's age: namely the time it has left to live (which affects the future value of any information it gains from investigating) and the time it has lived before encountering the novel object (which affects the agent's estimate of the object's abundance in the environment, and hence the future value of information it gains from investigating).

The primary aim of developing our model was to identify quantitatively the conditions under which the tendency of an agent to explore an unfamiliar object would change with age. Our secondary aim was to use the model to generate testable predictions. Thus, if the observed changes in exploratory tendency can be readily understood from an informational perspective, then the model may also make some additional predictions that can be can be evaluated by experiment. Third, in presenting our model we hope to further highlight the value of exploration–exploitation models (otherwise known as 'bandit models', Press, 2009) to represent and resolve the ubiquitous trade-off between gaining new information and exploiting current information. Greenberg and Mettke-Hofmann (2001) call for a cost–benefit approach to understand adaptive decision making when organisms encounter unfamiliar objects, and this is precisely what we have attempted to present here. Our model is intended to be general, but to help frame it, one may think of the dilemma facing a forager that encounters a novel prey type (or object) for the first time in its life and has to decide whether to investigate it (i.e. 'handle'), or leave it alone.

METHODS

Consider a simple system in which agents have a fixed time horizon of T to make decisions. This period may reflect the life span of the agent or the duration of a foraging season, so long as the individuals remain in the same stage of development (rats, for instance, exhibit developmentally dependent responses to novel stimuli; see Ricceri, Colozza, & Calamandrei, 2000) and the environment is unchanging. At age x ($0 \leq x \leq T$) in the agent's life, it comes across an object it has never seen before. The agent does not know how common these objects are in its environment, but it does know that this is the first such object it has encountered in its x time units of existence. So, if the agent encounters the object late (or early) in its life, then it is reasonable to infer that the object is rare (or common). As Greenberg and Mettke-Hofmann (2001) note, the primary purpose of exploration of novel objects is the acquisition of information, such as whether the object can be approached without consequence and whether the object provides a resource that can be used now and in the future. We therefore assume that the agent does not know the true probability that the object will be profitable (net benefit b after handling) or unprofitable (net cost c after handling) although it will have prior beliefs about the probability of the novel object being profitable to handle based, for example, on generalizing its previous experiences (Gershman & Niv, 2015). These net benefits and costs take into account the fact that investigating a novel object may take up valuable time and energy that could instead be invested elsewhere (i.e. opportunity costs).

Despite the agent's lack of understanding as to how common the object is and how profitable it is to handle, it still has to decide whether to investigate its properties (i.e. interact with it) or ignore it (move away)! Thus, it has to make a decision based on incomplete information. To identify the optimal behaviour for an agent, we first need to express the uncertainty in a quantitative way, and then identify a means for quantifying the future value of any