Animal Behaviour 82 (2011) 85–92

Contents lists available at ScienceDirect

Animal Behaviour

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

The evolution of superstition through optimal use of incomplete information

Kevin R. Abbott*, Thomas N. Sherratt

Department of Biology, Carleton University

ARTICLE INFO

Article history:
Received 16 November 2010
Initial acceptance 1 February 2011
Final acceptance 21 March 2011
Available online 8 May 2011
MS. number: A10-00872

Keywords:
Bayesian learning
causal learning
exploration–exploitation trade-off
optimization
superstition
two-armed bandit problem

While superstitions appear maladaptive, they may be the inevitable result of an adaptive causal learning mechanism that simultaneously reduces the risk of two types of errors: the error of failing to exploit an existing causal relationship and the error of trying to exploit a nonexistent causal relationship. An individual's exploration–exploitation strategy is a key component of managing this trade-off. In particular, on any given trial, the individual must decide whether to give the action that maximizes its expected fitness based on current information (exploit) or to give the action that provides the most information about the true nature of the causal relationship (explore). We present a version of this 'two-armed bandit' problem that allows us to identify the optimal exploration–exploitation strategy, and to determine how various parameters affect the probability that an individual will develop a superstition. We find that superstitions are more likely when the cost of the superstition is low relative to the perceived benefits, and when the individual's prior beliefs suggest that the superstition is true. Furthermore, we find that both the total number of learning trials available, and the nature of the individual's uncertainty affect the probability of superstition, but that the nature of these effects depends on the individual's prior beliefs.

> *… every man must judge for himself between conflicting vague probabilities.*
>
> **(Charles Darwin 1887, page 307)**

> *Dave knew perfectly well that making five copies [of the chain letter] and sending them to his friends wasn't going to bring him good luck. It was the bad luck he was worried about. … It's a hard world. You can't be too careful. It's not such a big deal to make five photocopies. Even at forty-six cents, a stamp is still a bargain.*
>
> **(Stuart McLean 2000, page 42)**

Superstitious behaviours can be defined as actions (or inactions) that are given in order to affect the probability that a beneficial outcome occurs when, in fact, there is no causal relationship between the action and the outcome (Skinner 1948). A stricter definition of a superstition (which, following Hood 2010, we henceforth refer to as a 'supernatural superstition') is one where there are no rational grounds to believe in a relationship between action and outcome, so that the agent's prior belief is that the relationship is unlikely. World leaders, athletes and media celebrities have all admitted to engaging in such superstitions (supernatural and otherwise), ranging from carrying 'lucky coins' to not treading on white lines (reviewed in: Wargo 2008; Hood 2010). Superstitions are challenging to explain from an evolutionary perspective when the action carries a cost, because in these instances one might expect that individuals engaging in superstitions would have lower fitness than individuals not engaging in these behaviours. Several authors have resolved the paradox by acknowledging that individuals that wish to exploit causal relationships in the environment must rely on incomplete information about causality, and that superstitious behaviours may be an unavoidable deleterious side effect of adaptively utilizing this information (Killeen 1978; Beck & Forstmeier 2007; Foster & Kokko 2009). This information, albeit incomplete, can be generated by natural selection (i.e. instinct) (Foster & Kokko 2009), cultural transmission, personal learning, or a combination of all three (McNamara et al. 2006; Beck & Forstmeier 2007). Here we focus primarily, although not exclusively, on personally learned information, not least because learning strategies may play a dominant role in superstition formation (Beck & Forstmeier 2007).

Skinner (1948) identified what he considered superstitious behaviour in pigeons, namely idiosyncratic behaviours (including head swinging or turning anticlockwise) in advance of food delivery, despite the fact that food was provided at fixed intervals. In so doing, he proposed a well-known learning-based explanation for these behaviours, suggesting that they might be the result of chance early pairings of certain actions and beneficial outcomes. For example, if an outcome occurs by chance

* Correspondence: K. R. Abbott Department of Biology, Carleton University, 1125 Colonel By Drive, Ottawa, ON K1S 5B6, Canada.
E-mail address: kabbott3@connect.carleton.ca (K.R. Abbott).

0003-3472/$38.00
doi:10.1016/j.anbehav.2011.04.002

after performing an action, then this might suggest to the agent that the action increases the probability that the outcome occurs. Even if there is no causal relationship between the action and outcome, such superstitions can be maintained if the agent ceases to explore the consequences of alternative actions, because there is too much to lose if the relationship turns out to be valid.

Nevertheless, despite Skinner's experimentally inspired hypothesis, and multiple formulations of the same basic idea (reviewed in Wargo 2008), the relationship between optimal exploration–exploitation strategies and the development of superstitions remains almost entirely unexplored from a quantitative perspective. Previous models that have attempted to elucidate the adaptive significance of superstitious behaviour (Beck & Forstmeier 2007; Foster & Kokko 2009) did not consider the learning process explicitly. While Beck & Forstmeier (2007) did raise the possibility that an adaptive learning strategy might be important, they did not explore the idea in detail. In this paper, we address this shortfall. Thus, we deal directly with how prior belief, chance events and the optimal use of incomplete information can combine to generate superstitious behaviours over multiple trials, rather than considering costs and benefits in one-off decisions (Foster & Kokko 2009).

Consider an individual that is attempting to determine whether a given action is causally related to a given outcome. What the individual needs to learn is whether the outcome is more or less likely to occur if the action is performed than if some other action is given (including nonaction). This requires exploring (i.e. testing) all possible actions. However, on most trials, the information the individual has already acquired will tentatively suggest that one particular action is associated with highest fitness returns. This can create a trade-off between giving the response that maximizes expected fitness on any given trial based on current knowledge (exploitation) and giving the response that provides the most information (exploration) about the true nature of the causal relationship between the individual's actions and outcomes (Cohen et al. 2007). As we will see, the payoff maximizing strategy should include a mechanism that allows an individual to explore in early trials and shift to exploitation in later trials.

MODEL DESCRIPTION AND RESULTS

In this paper, we present and analyse a model of optimal exploration–exploitation strategies for an individual learning the causal relationship between actions and outcomes. To make the model as tractable and informative as possible, we make several simplifying assumptions about the nature of proximate costs and benefits associated with actions and outcomes. We assume that there are only two possible mutually exclusive actions that the individual can perform on any given trial: one that is costly, A_c, and one that is cost-free, A_f. One of these actions can, as necessary, be considered a 'nonaction' (i.e. the behaviour of not adopting the alternative action). We also assume there are only two possible mutually exclusive outcomes to a trial: a relatively beneficial outcome, O_+, and an outcome, O_- that is less beneficial relative to O_+. For simplicity, we arbitrarily assume that the value of O_- is 0, the benefit of O_+ is $b = 1$, and the cost of A_c, c, is as some fraction of b.

There are two relevant conditional probabilities, $\Pr(O_+|A_c)$ and $\Pr(O_+|A_f)$. We assume that the individual knows the value of one of these conditional probabilities (the case where both values are known is trivial; the case where there is uncertainty associated with both values is interesting, but complex and is left for future papers). In the next subsection (The Costly Exploration Case), we develop the model and describe the results for the case where $\Pr(O_+|A_c)$ is unknown. We then briefly describe how the model can be modified for the case where $\Pr(O_+|A_f)$ is unknown (see The Cheap Exploration Case) and show the relevant results. Finally, we describe how the results differ between the two cases (see Costly versus Cheap Exploration).

The Costly Exploration Case

Model description

If A_f represents cost-free nonaction, then an individual will probably have had extensive experience with A_f, and should therefore know the probability $p_k = \Pr(O_+|A_f)$ with certainty (see Table 1 for a description of all mathematical terminology). At the same time, if the costly action, A_c, is relatively novel, then the individual will not know the exact size of the (hence unknown) probability $p_u = \Pr(O_+|A_c)$. We refer to this as the 'costly

Table 1
Description of model parameters and variables

Parameter/ variable	Description	Value/details	
		Costly exploration case (e.g. lucky charms)	Cheap exploration case (e.g. the number 13)
O_+	Relatively profitable outcome		
O_-	Relatively unprofitable outcome		
A_c	Relatively costly action	Costly novel action	Costly familiar action (e.g. actively avoiding A_f)
A_f	Relatively cost-free (cheap) action	Cheap familiar action (e.g. passively avoiding A_c)	Cheap novel action
b	Benefit of O_+ relative to O_-	1	Same in both cases
c	Cost of A_c relative to A_f	c is represented as some fraction of b	Same in both cases
p_k	True and known probability of getting the beneficial outcome, O_+, with the highly familiar action	$\Pr(O_+\|A_f)$	$\Pr(O_+\|A_c)$
p_u	True, but unknown, probability of getting the beneficial outcome, O_+, with the novel action	$\Pr(O_+\|A_c)$	$\Pr(O_+\|A_f)$
α, β	Parameters of the Beta prior p_u distribution that defines the individuals prior beliefs about the true value of p_u (≥ 1)		
N	Total number of trials available		
n	Current number of exploration trials ($\leq N$)	Number of A_c responses	Number of A_f responses
r	Current number of exploration trials on which the beneficial outcome, O_+, occurred ($\leq n$)		
$E_u(r,n)$	Expected value of the Beta distribution that defines the individuals current beliefs about the true value of p_u	$(\alpha+r)/(\alpha+\beta+n)$	Same in both cases
$F(r,n)$	The maximum expected return on all future trials	$\max((N-n)p_k, S(r,n)-c)$	$\max((N-n)(p_k-c), S(r,n))$
$S(r,n)$	The long-term expected return if the individual performs at least one more exploration trial	$E_u(r,n)(1+F(r+1, n+1))+(1-E_u(r,n))F(r,n+1)$	Same in both cases