



Finding the evolutionarily stable learning rule for frequency-dependent foraging

Steven Hamblin*, Luc-Alain Giraldeau

Département des Science Biologiques, Université du Québec à Montréal

ARTICLE INFO

Article history:

Received 28 November 2008
 Initial acceptance 29 January 2009
 Final acceptance 29 July 2009
 Available online 9 October 2009
 MS. number: A08-00764R

Keywords:

foraging
 learning rules
 producer–scrounger game

The cognitive mechanisms by which an organism comes to employ an optimal response to a situation are of great interest in behavioural ecology, but the basis and form of these mechanisms have been little studied. One approach employs learning rules, which are mathematical expressions that calculate the value of the behavioural alternatives in an organism's repertoire based on past and present rewards to those alternatives. Previous work on these learning rules has examined the performance of rules to determine whether they can achieve evolutionarily stable optimums. However, not only has this work tested rules in isolation, but the parameter values chosen to test them have been few and chosen arbitrarily. Moreover, the environments in which the rules have been evaluated are unchanging, a condition that does not favour learning. In this study we extend simulation work on three learning rules (relative payoff sum, linear operator and perfect memory). We use a genetic algorithm to both estimate the optimal parameter values for each rule and place the rules in competition with each other in a foraging game with a changing environment. Our results confirm earlier findings that the relative payoff sum is an ES learning rule. However, the results go much further because they show that the form of the learning rule that qualifies as evolutionarily stable combines near inextinguishable producing with highly responsive scrounging. The relative payoff sum may provide a single rule that can account for the way an animal's ecology can come to affect its specific set of learning sensitivities.

© 2009 The Association for the Study of Animal Behaviour. Published by Elsevier Ltd. All rights reserved.

In the course of their daily activities many animals are faced with choices. These may concern habitat selection, the exploitation of food patches, selecting prey to attack or mates to court or even opponents with whom to compete. The expected fitness return from such decisions depends on the information that is available to the animal and how it is used to decide on a course of action. When the value of alternatives is uncertain and changes over time, animals collect information or sample before deciding on a course of action. Selection is expected to have favoured individuals capable of acquiring information and acting appropriately (Stephens 1991). Models of optimal sampling under mostly nonsocial circumstances have been proposed and supported experimentally (Shettleworth et al. 1988).

In social situations sampling can become more difficult because the values of alternative courses of action not only depend on how the environment changes over time but also on the decisions of others (Giraldeau & Caraco 2000). In such cases, sampling is best analysed as an evolutionary game (Maynard Smith 1982). For example, an animal foraging for two prey types may face

uncertainty concerning the abundance of either prey type. But if it forages in the company of others, it must also contend with the changing abundances that will be induced by the prey selection decisions of its competitors. Determining which prey selection policy provides the greater payoff will require some sampling and adjustment of decision as conditions change.

The rule that governs an animal's sampling and its influence on decisions in the context of games has been labelled a learning rule (Harley 1981). Learning rules are mathematical descriptions of how animals assign values to behavioural alternatives based on current and past information about their payoffs obtained by sampling. Maynard Smith (1982) and Harley (1981) argued that natural selection would favour learning rules that led the group most quickly to the expected evolutionarily stable strategy: the ESS (Maynard Smith 1982). They defined an evolutionarily stable (ES) learning rule (Harley 1981) as a rule which, once fixed in a group, could not be invaded by any mutant rule. An ES learning rule, if one exists, would prove an extremely powerful tool for predicting the strategic use of alternative behaviour by animals engaged in games such as cooperation, fighting, habitat choice and resource exploitation (Dugatkin & Reeve 1997).

Maynard Smith (1982) and Harley (1981) proposed that one learning rule, the relative payoff sum (RPS), was a strong contender for the title of an ES learning rule. Not too surprisingly, therefore, a considerable amount of research has been focused on testing the

* Correspondence: S. Hamblin, Département des Science Biologiques, Université du Québec à Montréal, C. P. 8888, Succ. Centre-ville, Montréal, Québec H3C 3P8, Canada.

E-mail address: hamblin.steven@courrier.uqam.ca (S. Hamblin).

evolutionary stability of the RPS and other rules; some theoretical (Harley 1983, 1987; Hines & Bishop 1983; Houston 1983; Houston & Sumida 1987; Tracy & Seaman 1995), some empirical (e.g. Milinski 1984; Kacelnik & Krebs 1985; Amano et al. 2006), and some using computer simulation (e.g. Regelmann 1984; Beauchamp 2000, 2004; Beauchamp & Fernández-Juricic 2004; Beggs 2005; Amano et al. 2006; Spataro & Bernstein 2007). In the course of this research, three rules have been proposed as serious contenders to the title of ES rule: the relative payoff sum (RPS; Harley 1981), the linear operator (LOP; Bush & Mosteller 1955), and perfect memory (PM; Houston & Sumida 1987), the mathematical details of which are discussed in the *Methods*.

Despite the considerable research effort devoted to ES learning rules, none has emerged as convincingly evolutionarily stable. Strong conclusions have been hampered by a combination of three major problems: first, few studies have pitted all rules against all others; second, all studies have used rules fitted with just a few parameter values, apparently chosen arbitrarily; third, almost all rules have been tested in unchanging environments where learning is of little value.

One of the most complete investigations into the question of ES learning rules has been conducted by Beauchamp (2000). To our knowledge, this is the only study that has attempted to study ES properties of learning rules in competition against each other. To do so, Beauchamp used an agent-based model but tested rules using only three parameter values that were apparently chosen arbitrarily. Harley's (1981) original concern of whether a learning rule might lead a population to adopt an uninvadable strategy through behavioural means cannot be addressed without determining the parameters that drive these learning rules. To date, all studies of the ES properties of learning rules have neglected the importance of the exact parameters used when testing the rule. Given the possible set of parameter values for these rules, only a minuscule portion of the available parameter space has been explored. If we are to conclude about a rule's evolutionary stability, it is imperative that we do so knowing that a rule's success is not just due to its being fitted with better performing parameters. The only way to do this is by using each rule's optimal set of parameters.

How do we find optimal parameter values for these rules to test them with? An analytical solution to these equations is out of reach, and the set of possible parameter values (the parameter space) is too large to feasibly conduct an exhaustive search. To solve this problem, we turn to a heuristic search technique known as a genetic algorithm. Genetic algorithms mimic the techniques of natural selection (differential reproduction, mutation and recombination) to find solutions to optimization problems (Sumida et al. 1990; Huse et al. 1999). Here, we use the genetic algorithm to evolve candidate rules with different parameter values, selecting those that perform best in the social foraging task to reproduce, and using mutation and recombination to sweep the parameter space. One of the virtues of the genetic algorithm approach to modelling is that the evolution of traits can be made explicit (Hamblin & Hurd 2007; Ruxton & Beauchamp 2008), and by linking it to a model with an explicit description of a behavioural mechanism, insights into the interaction between the two may be gained (reviewed in Seth 2007). A cautionary note is required here, though, since the language of the field of evolutionary computation borrows heavily from biology while the comparisons between the two are not always so clear. It is important not to take the jargon of genetic algorithms too literally, as the details of the genetic algorithm may seem odd, biologically; parameters and methods for genetic algorithms are usually chosen for search power, not biological realism (for a discussion on the ways of conceptualizing genetic algorithms in behavioural ecology, see Ruxton & Beauchamp 2008).

Beauchamp tested the properties of these rules as agents were engaged in playing one of two foraging games, the producer–scrounger game (Barnard & Sibly 1981) or an ideal free distribution game (Fretwell & Lucas 1969; Sutherland 1983). Given that Beauchamp's results were similar whether the agents were engaged in a PS or an IFD game, we chose to have our agents play only a PS game. Although most experimental work on PS games concerns foraging in small flocks of birds (Giraldeau & Caraco 2000), it is important to realize that the structure of the PS game is more general and applies widely to any case where the investment of some individuals is exploited by others (Barnard 1984), much like the caller and silent satellite male strategies of many toads (e.g. Howard 1978) or the digging versus entering strategies of digger wasps (Brockmann et al. 1979).

As in all investigations of the ES property of learning rules, Beauchamp's (2000) study is set in an unchanging environment. This is problematic because learning may be of most value when the environment is at an intermediate level of variability (Stephens 1991; Kerr & Feldman 2003). So testing a learning rule under conditions where learning is of little value may not have provided the rules with conditions that allowed them to perform efficiently. To provide rules with a variety of conditions, we focus on changing two variables: group size and environmental variability. Population size is relevant for both biological and technical reasons. Technically, genetic algorithms are more powerful with larger group sizes, since a large group (population) size means more candidate rules to select among. Biologically, we investigate the effect of group size to determine how the use of learning may change in large or small groups. With regard to environmental variability, we will investigate rule performance as environmental features such as patch density and food items per patch.

Our first objective is to go beyond the use of arbitrary parameters when testing rules. To do this, we determine the optimal parameters for each rule using a genetic algorithm. As this is occurring, the rules simultaneously compete with each other in the genetic algorithm, allowing us to form a clearer picture of the evolutionary stability of the rules, which is our second objective (Houston & Sumida 1987). Third, we wish to establish the evolutionary stability of rules in environments that vary. Finally, we hope to determine whether the rules predict similar or distinct group structures; that is, whether agents will end up specializing on one or the other alternative or instead become generalist individuals that switch from one to the other alternative.

METHODS

The Model

We use an agent-based foraging model that builds on the model described in Beauchamp (2000) and extends it by using a genetic algorithm to optimize the rule type and rule parameter choices of the agents in the model.

The Learning Rules

Each rule has its own peculiarities. The RPS has two components: one concerns the estimated value of the alternatives, the other is a decision based on these values.

Relative payoff sum (RPS):

$$S_i(t) = x \times S_i(t-1) + (1-x) \times n + P_i(t)$$

where $S_i(t)$ is the value placed on behavioural alternative i at time t , x is the memory factor that determines how highly the past is valued, n is the residual, the cutoff below which the valuation of the

Download English Version:

<https://daneshyari.com/en/article/2417545>

Download Persian Version:

<https://daneshyari.com/article/2417545>

[Daneshyari.com](https://daneshyari.com)