Published in final edited form as:

J Theor Biol. 2010 December 21; 267(4): 573–581. doi:10.1016/j.jtbi.2010.09.026.

Co-evolution of learning complexity and social foraging strategies

Michal Arbilly^{a,*}, Uzi Motro^b, Marcus W. Feldman^c, and Arnon Lotem^a

^aDepartment of Zoology, Faculty of Life Sciences, Tel-Aviv University, Tel-Aviv 69978, Israel

^bDepartment of Evolution, Systematics and Ecology and Department of Statistics, The Hebrew University of Jerusalem, Jerusalem 91904, Israel

^cDepartment of Biology, Stanford University, Stanford, CA 94305, USA

Abstract

Variation in learning abilities within populations suggests that complex learning may not necessarily be more adaptive than simple learning. Yet, the high cost of complex learning cannot fully explain this variation without some understanding of why complex learning is too costly for some individuals but not for others. Here we propose that different social foraging strategies can favor different learning strategies (that learn the environment with high or low resolution), thereby maintaining variable learning abilities within populations. Using a genetic algorithm in an agent-based evolutionary simulation of a social foraging game (the producer-scrounger game) we demonstrate how an association evolves between a strategy based on independent search for food (playing a producer) and a complex (high resolution) learning rule, while a strategy that combines independent search and following others (playing a scrounger) evolves an association with a simple (low resolution) learning rule. The reason for these associations is that for complex learning to have an advantage, a large number of learning steps, normally not achieved by scroungers, is necessary. These results offer a general explanation for persistent variation in cognitive abilities that is based on co-evolution of learning rules and social foraging strategies.

Keywords

Foraging; Cognition; Game theory

1. Introduction

Learning is a process of acquiring information, storing it in memory, and using it to modify future behaviors. It is believed to be adaptive because under a wide range of conditions it allows the learner to generate predictions about its environment, and hence to make better decisions, than by using innate knowledge alone (Johnston, 1982; Stephens, 1991; Bergman and Feldman, 1995). It would seem that survival is improved when learning abilities allow more details and more accurate information to be acquired from the physical environment; therefore we would expect natural selection to favor complex learning abilities, allowing

© 2010 Elsevier Ltd. All rights reserved.

*Corresponding author: Tel.: +972-3-6408381; fax: +972-3-6409403. michala8@post.tau.ac.il (M. Arbilly).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

high precision, over simple and less accurate ones. But variation in learning abilities is maintained in nature within species, from insects (Mery and Kawecki, 2002; Raine and Chittka, 2008) to mammals (Tolman, 1924; Tryon, 1940. See also Bouchard and McGue, 1981; Plomin, 2001). This variation suggests that complex learning may be costly and, at least under some conditions, does not provide the learner with the expected advantage over simpler learning rules.

Several models for the evolution of learning have included a cost for individual learning (e.g. Boyd and Richerson, 1988; Rogers, 1988; Stephens, 1991; Bergman and Feldman, 1995; Feldman et al., 1996; Wakano et al., 2004; Aoki et al., 2005). In recent years, experimental evidence for such a cost has been provided in *Drosophila melanogaster* (Mery and Kawecki, 2003; Mery and Kawecki, 2004; Burger et al., 2008) and in the parasitoid wasp *Biosteres arisanus* (Dukas and Duan, 2000). This cost may involve the energy required for the creation and maintenance of brain tissue (Johnston, 1982; Dukas, 1999), or the attention diverted from other cognitive tasks, such as scanning for predators or searching for food while exploiting existing knowledge (Dukas, 1998). Another cost of learning may be due to errors during the learning process (Johnston, 1982). Such errors may result in learning to recognize the wrong stimuli rather than the correct one (e.g., Boyd and Richerson, 1988; Lotem, 1993; Rodriguez-Girones and Lotem, 1999) or in sub-optimal decisions due to sampling errors and their dynamic consequences (e.g., Stephens, 1987; March, 1996; Bateson and Kacelnik, 1997; Niv et al., 2002; Shafir et al., 2008). All of these possible costs of learning are likely to increase with learning complexity, which may explain why complex learning may be more difficult to evolve unless it can be compensated by greater benefits. Yet high costs of complex learning cannot explain variation in learning complexity within populations, that is, it cannot explain why complex learning is too costly for some individuals but not for others, and how such variation in the costs and benefits of learning can be maintained over evolutionary time.

In this paper we explore the possibility that in social animals, variation in the ability to learn where to find food or other resources may be a result of co-evolution of individual learning rules and social foraging strategies. Social animals are frequently described as being engaged in an evolutionary game between "producers" that search for food independently, and "scroungers" that tend to follow other individuals and join the formers' food discoveries (Barnard and Sibly, 1981; Giraldeau and Caraco, 2000; Giraldeau and Dubois, 2008). Potential implications of these strategies for individual learning may include the following: first, individuals that rely on scrounging may be less dependent on their individual learning abilities. Second, in comparison to producers, scroungers will have fewer opportunities for individual learning because they spend part of their time following others. Thus, it would appear that the producer/scrounger tendencies create different conditions for the evolution of individual learning. Put simply, one can expect that in having more time to learn and more to gain from learning, producers should be "smarter" than scroungers. More precisely, we predict that if complex learning rules (that acquire information at a high resolution, i.e., a high degree of detail) result in better outcomes than simple learning rules (that acquire information at a low resolution) but at a cost of more learning steps, producers may be selected to use complex learning rules that scroungers cannot afford. Interestingly, this prediction has not been tested theoretically or empirically. There has been extensive research on the producer-scrounger game (reviewed by Giraldeau and Caraco, 2000; Giraldeau and Dubois, 2008), on the evolution of social versus individual learning (Boyd and Richerson, 1985; Boyd and Richerson, 1988; Rogers, 1988; Boyd and Richerson, 1995; Feldman et al., 1996; Wakano et al., 2004; Aoki et al., 2005; Borenstein et al., 2008), and on evolutionarily stable learning rules for choosing among strategies in a game (Harley, 1981; Houston and Sumida, 1987; Tracy and Seaman, 1995; Beauchamp, 2000; Hamblin and Giraldeau, 2009). There is also increasing interest in the evolution of individual learning in stochastic and

stable environments (Stephens, 1991; Bergman and Feldman, 1995; Kerr and Feldman, 2003; Groß et al., 2008). However, despite this rich background, the possible co-evolution of individual learning abilities and social foraging strategies has not been studied.

There is empirical evidence that the process of scrounging may block the learning of foraging cues in some species (Giraldeau and Lefebvre, 1987; Beauchamp and Kacelnik, 1991; Lefebvre and Helder, 1997), but not in others (Fritz and Kotrschal, 1999; Caldwell and Whiten, 2003). This suggests that at least in some cases, following others and relying on social information may interfere with learning individually (see also Laland, 2004; Galef and Laland, 2005). While these ideas touch on the interaction (and possible conflict) between individual and social learning mechanisms, the more basic problem we address here emerges even when scroungers do not use any form of social learning. It emerges as soon as scroungers spend more time scrounging from others and less time searching by themselves. Thus, our goal in this paper is to construct a theoretical platform that will be used to study the basic co-evolution of individual learning rules (rules to study the environment) and social foraging strategies (i.e. producing versus scrounging) before adding social learning rules or rules for learning to choose among social foraging strategies. Using a genetic algorithm in an agent-based evolutionary simulation of the producer-scrounger game, we test whether the tendency to search for food independently (i.e., to play a producer) evolves an association with a complex learning rule, while a partial tendency to follow others (i.e., to play a scrounger) evolves an association with a simple learning rule. Our results provide a theoretical explanation for persistent variation in cognitive abilities within populations that is based on association between learning rules and social foraging strategies.

2. The model

Our model for individual learning rules and social foraging strategies was inspired by observations on the house sparrow, *Passer domesticus*, in our research group (see Katsnelson et al., 2008), as well as by previous work on social foraging in this species (e.g. Barnard and Sibly, 1981; Liker and Barta, 2002; Lendvai et al., 2004; Liker and Bokony, 2009). Nevertheless, it may be applied to all social foragers.

2.1. The population

We simulated a population of n haploid agents. Each agent is defined by two genetic loci: a foraging strategy locus, F , which defines the agent's probability of applying the producer or scrounger strategy, and a learning rule locus, L , which determines the resolution at which the agent learns the environment when applying the producer strategy. All agents have the same lifespan, which is composed of J foraging steps executed simultaneously by all agents during a "producer-scrounger game" (see below). All agents end their life at the end of the game and there is no overlap between generations. Fitness is determined according to the agents' accumulated payoff during the J steps of the game, and the top 50% of agents are then selected to reproduce asexually, each giving rise to two genetically identical offspring (for an extension of this type of model to sexual reproduction see Arbilly et al., *Submitted*). Mutation rate μ for each locus is $1/n$, and the population is followed through G generations. The simulated environment does not change (see below), but because information about the environment is not transferred from one generation to the next, each generation experiences the environment as new. We can therefore assume that although individuals are engaged in a single game of J steps during their lifetimes, similar results would be expected for repeated games of J steps during each lifetime, as long as learning starts anew for each new game (i.e., repeated games are independent). The assumption that learning starts anew every generation may be realistic when the relevant cues for finding food change over time or when the environment changes.

2.2. The environment

The simulated environment consists of four food patches E_1 , E_2 , E_3 , and E_4 . The four patches are visually distinguishable from one another, but patches E_1 and E_2 have a common visual characteristic, for example, the same color, and so do patches E_3 and E_4 ; therefore, these can be viewed as E_{12} and E_{34} , respectively (see illustration in Table 1). Each patch is represented by two parameters: the probability of finding food items in the patch, and the nutritional value of these items (see Table 1). The two parameters are set to be negatively correlated: while the patch that always provides a payoff has the lowest expected value, the patch in which retrieving a payoff is a rare event is the patch with the highest expected value. Simply put, the best food (the food with the highest nutritional value) is the least likely to be found. This environmental setup represents a non trivial learning task and is likely to be common in nature (valuable food items may be depleted sooner and become less common). To keep the simulation sufficiently simple, patch parameters remained constant throughout the game (no depletion during the game), and the environment had no spatial structure.

2.3. Foraging strategy

The foraging strategy locus (F) determines the agent's probability of applying a producer strategy at each step of the producer-scrounger game described below (with the complementary probability of applying a scrounger strategy). In this model there are two alleles: F_1 , whose carriers have a probability 1.0 of adopting a producer strategy, and F_2 , whose carriers have a probability of only 0.5 of applying the producer strategy; otherwise they apply the scrounger strategy (i.e., to follow others). Thus carriers of the F_2 allele can also learn which food patch is most profitable and use this information in future steps, but they do so with probability one-half in each of the foraging steps (i.e., when they do not scrounge from others). The two producing probabilities defined by the two alleles (i.e. 1.0 and 0.5) were chosen in order to simulate a situation in which individuals differ considerably in their social foraging strategies but are still capable of using individual learning. Preliminary simulations have shown that bimodalities of this sort evolve as a result of the game even when multiple alleles coding for various probabilities are available in the simulation. The F_1 and F_2 alleles were assigned randomly with equal probability to the n individuals in the first generation.

2.4. Learning rules

Learning takes place only when an agent is applying the producer strategy (as explained in the Introduction, the evolution of social learning in this system is beyond the scope of the present paper). The value of the food item it has obtained (or zero in the case it has obtained no food) updates the agent's memory, and this information is used in the agent's subsequent producing steps to decide in which patch to search for food. As a typical memory updating rule we used the linear operator rule (McNamara and Houston, 1987; Bernstein et al., 1988; Beauchamp, 2000; Groß et al., 2008; Eliassen et al., 2009), which is also known as the weighted return rule (March, 1996). It yields for each patch a weighted average of the most recently acquired payoff and previous payoffs, according to the equations:

$$M_{i,t+1} = \alpha M_{i,t} + (1 - \alpha) Y_{i,t} \quad t=1, 2, \dots, J$$

where $M_{i,t}$ is the value of patch E_i in the agent's memory at step t , and $Y_{i,t}$ is the updating vector: $Y_{i,t}$ equals the payoff obtained at step t if patch E_i was visited at step t . Here α ($0 \leq \alpha \leq 1$) is a memory factor describing the weight given to the agent's past experience as the most recent experience updates memory. Note that Y_t is the payoff obtained by the producer when alone at the food source, and it may be different from the payoff the producer

eventually receives if joined by a scrounger (see game description below). For the first step, $M_{i,1}$ are set to be the mean of all four patches' expected values ($M_{i,1} = 0.53$ for all i).

Two learning rules were defined by the alleles at the learning rule locus. The first is a complex learning rule (allele L_1) that views the environment in full detail, as composed of the four food patches E_1, E_2, E_3, E_4 . The second is a simple learning rule (allele L_2) that views the environment as composed of only two patches, E_{12} and E_{34} . Note that with the complex learning rule individuals can potentially learn to prefer the patch with the highest expected value (Table 1, patch E_1 , expected value = 1.0), while with the simple learning rule they can at best learn to prefer the E_{12} combination (with an expected value of 0.75). Thus, if applied successfully, complex learning should be better than simple learning. However, since complex learning divides its sampling steps between four patches in memory representation (rather than two) it is more likely to involve sampling errors when the total number of learning steps is small. A third allele L_0 results in no learning, and for carriers of L_0 patch information is not stored in memory and the patch in which to forage when producing is chosen at random (with an expected payoff of $M_{Ei,t} = 0.53$ remaining constant throughout the game). In the first generation all agents are assigned allele L_0 , and alleles L_1 and L_2 are introduced into the population by random mutations.

To account for the possibility that complex learning may be more costly than simple learning (see introduction), a cost measured by the parameter γ was incurred by individuals carrying the complex learning allele (L_1). This cost was a fractional reduction by γ of the final cumulative payoff (after the j steps) to agents carrying the complex learning rule allele, L_1 .

2.5. Decision rule

When producing, the agents decide in which patch to forage based on the information acquired and updated in the abovementioned learning process. We used the exponential (logit) response rule (following Ben-Akiva and Lerman, 1985; Busemeyer and Myung, 1992; Bereby-Meyer and Erev, 1998; Camerer and Ho, 1999), according to which patch E_i is chosen with a probability based on its relative weight in memory (M_{Ei}), using the expression:

$$p_{i,t+1} = \frac{e^{M_{i,t}}}{\sum_{k=1}^4 e^{M_{k,t}}}.$$

Here $p_{i,t+1}$ is the probability of choosing patch E_i at step $t+1$. The above expression holds for complex learners (L_1 carriers) that distinguish between all four patches; simple learners (L_2 carriers) use the same expression but k goes from 1 to 2 (rather than from 1 to 4). The use of an exponential ratio prevents choice rigidity and allows more frequent sampling of seemingly less attractive patches than with a simple ratio.

2.6. The producer-scrounger game

As described above, the agents' lifespan is composed of J steps. Each step begins with a choice of foraging strategy (producing or scrounging) according to the agent's foraging allele (see above). Agents with the F_1 allele (they always search), or agents with the F_2 allele (they use the producer strategy with a probability of 0.5), choose a patch to forage in (based on the above learning and decision rules), and receive a payoff (or not) according to patch parameters described in Table 1. Agents with the F_2 allele that act as scroungers are then assigned to producers to simulate a scrounging process. We assume that at each step

scroungers are able to distinguish successful producers from producers that failed to find food. A producer in our game can be joined by only one scrounger, and the payoff is divided between the two agents (such that a finder's share of half the payoff always goes to the searcher). This assumption is different from some other producer-scrounger models where a producer can be joined by many scroungers (see for example Beauchamp, 2000; Hamblin and Giraldeau, 2009). We chose to restrict the number of joiners to one because our observations in flocks of sparrows indicated that small food patches cannot usually accommodate more than two sparrows at a time. It is important to note, however, that this assumption does not change the negatively frequency-dependent dynamics of the game. Since assignment of scroungers to producers is random, some producers may not be joined by any scrounger (especially when scrounging is rare), while it is possible for a scrounger to be left with no payoff if the producer it was assigned to has already been joined by another scrounger. This probability is obviously greater when the proportion of scrounging increases. Producers, on the other hand, still secure their finder's share and can do relatively better than scroungers when scroungers become too common (see also Sibly, 1984). This type of model is not affected by group size (unless the size is very small); however, this parameter is irrelevant to our analysis as we do not seek to determine the ESS.

Simulation parameter symbols and meanings are summarized in Table 2.

3. Results

3.1. The effect of number of steps on the success of simple and complex learning rules

We first tested our initial prediction that complex learning would be better than simple learning only after a sufficient number of learning steps. To that end we assessed the effect of the three learning alleles on the cumulative payoff of pure producers (agents carrying F_1) after J learning steps by simulating a single generation for 100,000 populations of $n = 100$ individuals (with the three learning alleles assigned with equal probability) and by monitoring in each population the genetic profile of the 50% of individuals with the highest payoff after J steps (where $5 \leq J \leq 100$) during the simulation. Note that this process does not represent an evolutionary trajectory but produces the fitness comparisons for a single generation. The results (Fig. 1a) show that each learning allele had an advantage for a different range of number of steps. When the memory factor α was set to 0.5, with a small number of steps ($J < 14$) learning was in fact maladaptive and the non learning allele L_0 prevailed; in the medium range ($14 \leq J \leq 27$), the simple learning rule (allele L_2) was the best strategy, and for $J > 27$, the complex learning rule (allele L_1) started to prevail, and its relative advantage increased with the number of steps. Qualitatively similar results were obtained with different values of the memory factor α , such as $\alpha = 0.2$ or 0.8 . Thus, the results support our initial prediction that complex learners can obtain better payoffs than simple learners, but only after a sufficiently large number of steps. Interestingly, the results also show that when the number of learning steps is very small ($J < 14$), the random patch choice of the non learning allele (L_0) provides higher payoffs than either the simple and complex learning rules. This counterintuitive result can actually be explained by the tendency of such learning algorithms to develop risk aversion (March, 1996; Niv et al., 2002, see also Dornell and March, 2001 for the "hot stove phenomena"). That is, initial visits to patches E_1 or E_2 may result in a zero payoff despite their high expected value (see Table 1), and thus they will be avoided during the next few steps. With a small number of steps, there will not be enough opportunities to correct this erroneous choice (an initial preference of E_3 and E_4), which causes many of the learners to do worse than the non learners.

To confirm the above results for the context of the producer-scrounger game we repeated the above analysis, but this time with populations that carry the F_2 allele (part-time scroungers), thus allowing a producer-scrounger game to emerge between the mixed players (part-time

scroungers) of the tested generation. The results (Fig. 1b) are generally similar to those obtained with pure producers, but the differences in cumulative payoffs between the three learning alleles were smaller (compared with Fig. 1a). Also, considering that part-time scroungers had only about half of the steps available for learning, the similar transition point from simple to complex learning (around 30 steps) indicates that under game conditions, complex learning can be advantageous after fewer learning steps. The differences between the two analyses can be attributed to game dynamics (see discussion) and to the fact that payoffs are also determined by the benefits accumulated while playing a scrounger (which are not affected by the learning allele). In both cases, however, there is an intermediate number of steps during which simple learning (the L_2 allele) is better than complex learning (the L_1 allele).

It should be noted that while truncation selection of the top 50% of the population may appear as strong selection, the small differences in survival of the learning alleles, shown in Fig. 1, indicate that the selection acting in our simulation is rather weak.

3.2. Evolution of simple and complex learning rules in a producer-scrounger game

Simulation results for a population of $n = 300$, allowed to evolve for 3,000 generations, with various values of J (number of steps in a game) and γ (cost of complex learning) are presented in Fig. 2. For all parameter sets, the non-learning allele L_0 , which was the only allele at the learning locus in the first generation, rapidly decreases in frequency and comes close to extinction, as the other two alleles – simple learning and/or complex learning – take over the population. For $\gamma = 0$, there is still a cost for complex learning but this is only the cost inherent to the longer learning process of this algorithm that requires more learning steps to become profitable (see above). Yet, despite this cost, both foraging strategies adopt the complex learning rule (Fig. 2a, e, i, m). An example of evolutionary dynamics of one simulation for $G = 3000$ generation, $J = 50$ steps, $\gamma = 0$, is shown in Fig. 3a. The success of the complex learning rule while also being used by part-time scroungers (in a game of 50 steps) appear to contradict our initial results (Fig. 1); this will be clarified later.

As we increase γ , a pattern of association between foraging strategy and learning rule emerges: As might be expected, higher cost of complex learning (γ) makes complex learning less likely to evolve, while an increase in the number of steps (J) favors its evolution. Thus, for combinations of a high cost and a relatively low number of steps, both foraging strategies evolve an association with simple learning (Fig. 2b, c, d, h, l), while for combinations of a low cost and a relatively high number of steps, both foraging strategies evolve an association with complex learning (Fig. 2e, f, i, j, m, n, o). We can display these conclusions more precisely by considering the relative difference between the frequency of

the complex and simple learners among the pure producers (ΔF_1 , defined as $\frac{\#F_1L_1 - \#F_1L_2}{\#F_1L_1 + \#F_1L_2}$, where $\#F_1L_1$ is the number of individuals carrying the F_1L_1 haplotype, etc.), and among the

part-time scroungers ($\Delta F_2 = \frac{\#F_2L_1 - \#F_2L_2}{\#F_2L_1 + \#F_2L_2}$). In order to evaluate the effect of learning cost (γ) and the number of steps (J) on ΔF_1 , we performed a two-way ANOVA. The cost has a significant and decreasing effect on ΔF_1 ($F_{3,144} = 360.27$, $P < 10^{-6}$); the number of steps has a significant and increasing effect on ΔF_1 ($F_{3,144} = 266.90$, $P < 10^{-6}$); and the interaction is also significant ($F_{9,144} = 28.14$, $P < 10^{-6}$). Likewise for ΔF_2 ($F_{3,144} = 477.03$, $P < 10^{-6}$; $F_{3,144} = 90.33$, $P < 10^{-6}$; $F_{3,144} = 16.49$, $P < 10^{-6}$, respectively). Both ΔF_1 and ΔF_2 are decreasing with γ . For small enough values of γ they are both positive, and for large enough values of γ they are both negative. However, there is an intermediate range of γ where ΔF_1 is positive, while ΔF_2 is negative (Fig. 4). In this intermediate range, the pure producer strategy F_1 is associated with the complex learning rule L_1 , and the part-time

scrounger strategy F_2 is associated with the simple learning rule L_2 (Fig. 2g, k, p). An example of the evolutionary dynamics at $J = 50$, $\gamma = 0.01$, is shown in Fig. 3b. Under these conditions, the simple and complex learning rules coexist in the population, each in association with a different social foraging strategy.

Finally, to confirm that our results did not reflect premature convergence stemming from the parameters of the genetic algorithm, we also ran the simulations with lower selection thresholds (eliminating only the lowest 20% or 40% of agents from the population rather than the lowest 50%) and with an increased mutation rate ($\mu = 10/n$ for each gene). We found no evidence for a tendency to converge to a solution different from the one represented by our main results; as expected, lowering the selection threshold and increasing mutation rate prevented complete removal of the inferior haplotypes from the population and was more likely to result in highly polymorphic populations (not shown).

4. Discussion

Using an agent-based evolutionary simulation we were able to demonstrate a situation in which complex and simple learning rules evolve stable associations with alternative social foraging strategies in the producer-scrounger game. In the following, we will attempt to explain these results, and discuss their implications for the possible co-evolution of learning rules and social foraging strategies in natural populations.

4.1. Learning rules and the number of learning steps

Our initial non-evolutionary analysis explored the relationship between the number of learning steps available for the learner and the success of simple and complex learning rules. We found that a minimum number of learning steps is necessary even for the simple learning rule to be successful, and that a larger number of learning steps was required for the success of the complex learning rule. As explained earlier, this is not merely the result of a large statistical error due to the small sample, but is mainly a result of the highly variable environment that gives rise to risk aversion (March, 1996; Derrnelli and March, 2001; Niv et al., 2002). Interestingly, our analysis suggests that the effect of this bias on a population of learners decreases if learners have a sufficiently large number of learning steps (and as long as the initial values given to patches in memory are greater than zero). In this case, some learners will eventually revisit the better patches (with a probability dictated by the decision rule in our model) and might be lucky enough to experience the high reward (e.g. 4 and 1.5 in our E_1 and E_2 patches). It is probably enough that this happens to only some of the learners in order to make learning better on average than non-learning, and, with additional learning steps, to make complex learning better on average than simple learning. The present analysis may not be the only situation under which the number of learning steps affects the relative success of simple and complex learning rules. However, it provides an explicit example of how the number of learning steps can limit the number of cues that can be learned effectively when searching for food in a variable environment.

4.2. Learning rules, the number of steps, and game dynamics

In our second analysis, we tested the relationship between the number of learning steps and the success of simple and complex learning rules under the dynamic conditions of the producer-scrounger game (see Fig. 1b and related text). Although the results were generally similar, we note two differences. First, the differences in payoffs between learning strategies were smaller. As mentioned earlier, this may be explained by the payoffs that can be gained while playing a scrounger in about half of the steps, which are similar for the two types of learners. The second and more surprising difference is that under game conditions, complex learning needed fewer learning steps to become better than simple learning. This was

indicated by the transition point between the two learning rules in Fig. 1a and 1b that remains the same despite the fact that under game conditions (Fig. 1b), the players (all of which were part-time scroungers) had only about half of the steps available for learning. This positive effect of game conditions on the success of complex learning was also confirmed by our evolutionary simulations: In contrast to Fig. 1a, where complex learning was best only after 27 learning steps, in the evolutionary simulations, as long as the additional cost for complex learning was zero or very small, part-time scroungers adopted complex learning even in games of 30 or 50 steps, when the number of steps at which they learned was probably less than 27 (Fig. 2a, 2e, and 2f). That game dynamics assisted complex learning illustrates why explicit simulations of learning processes and game dynamics are important. Under game conditions, the payoff accumulated by part time scroungers is determined by three factors: the number of steps required to learn to search for food in the best patch (and the resulting payoff), the payoff lost to scroungers while producing, and the payoff obtained while scrounging. Under game conditions, when part-time scroungers apply each strategy with equal probability, the actual number of learning steps available to each individual varies as a result of this coin-flipping strategy choice. Interestingly, most of the improved success of complex learning under game conditions (see above) disappears when this stochasticity is removed in a simulation of a deterministic strategy choice with alternating strategies after each step (not shown). The reason that variation in the number and sequence of learning steps improves the success of complex learning is that it allows a slightly larger number of complex learners to find the highly rewarding food in the risky patches, and to learn to prefer them for the rest of the game. For similar reasons, the payoff producers lose to scroungers while producing, and the payoff obtained while scrounging also favor complex learning under game conditions but to a smaller extent (unpublished analysis, available upon request).

4.3. Co-evolution of learning and social foraging strategies: should producers be "smarter" than scroungers?

Our evolutionary analysis clearly identified a range of conditions under which producers evolve an association with a complex learning rule, while part-time scroungers evolve an association with a simpler learning rule. This result emerged from a combination of a very small cost of 1% for complex learning, and a limited number of steps in the game. It is interesting to note that in the absence of cost for complex learning, the lower number of learning steps available for part time scroungers was insufficient to create this dichotomy. However, by adding a slight cost for complex learning, the dependence of learning on of the number of learning steps became critical, allowing only pure producers to evolve strong association with complex learning.

Are the conditions producing our theoretical results likely to be common in natural populations? It is clear from our analysis that when the number of steps is sufficiently large, both producers and part-time scroungers can use complex learning, and in this case alternative social foraging strategies cannot explain variation in learning abilities within populations. We suggest, however, that such conditions that favor "complex learning for all" represent mainly the constraints of our simulation, and that realistically, the transitional zone where each social foraging strategy favors a different learning rule will be expanded. In our simulation we considered only four types of patches and two types of learning rules. In the real world, however, it is easy to see that when part-time scroungers can afford a learning rule of a given complexity, pure producers should almost always be able to use an even more complex learning rule. In other words, the logical principle illustrated in our simulation for the case of two versus four patch types in memory representation, should also hold for three versus five, or for six versus ten. As long as social foraging strategies differ

substantially in the number of learning steps available for the players, a differential effect on the complexity of learning is almost inevitable.

Our current analysis ignores the problem of genetic recombination by making the simplifying assumption of asexual reproduction. In our model, the haplotypes representing the successful combinations of genes (i.e. F_1L_1 and F_2L_2) are always passed to the next generation as exact copies of themselves. This is clearly unrealistic for most organisms that are engaged in producer-scrunner games since they reproduce sexually (e.g. birds and mammals). In such organisms, inferior combinations of alleles can emerge each generation as a result of genetic recombination. We address this problem thoroughly in a separate paper (Arbilly et al., *Submitted*), where we show that the predictions of our analysis can be generalized.

The possible co-evolution of genes for certain behaviors and genes for learning ability has been recently demonstrated in *Drosophila* by Mery et al., (2007). They show that variation in a gene associated with foraging behavior also affects learning abilities: larvae which are homozygous for one variant and tend to move more while feeding also exhibit better short-term learning and poorer long-term memory than larvae that are homozygous for the second variant and are less active feeders. Although it is not clear what stabilizes the two behavioral morphs in this case, the association of each of them with a different type of learning is quite consistent with the spirit of our results. Another relevant body of literature that deals with the coupling of two or more behavioral traits in the course of evolution is the rapidly growing field of behavioral syndromes, or animal personalities (reviewed by Sih et al., 2004a; Sih et al., 2004b; van Oers et al., 2005; Bell, 2007; Sih and Bell, 2008). While there is increasing empirical evidence for behavioral polymorphism that is based on sets of correlated traits, theoretical models explaining the evolution of such behavioral syndromes are just starting to emerge (Wolf et al., 2007; Wolf et al., 2008). Similar to our approach, these models consider the critical role of frequency-dependent selection in stabilizing different morphs, and the importance of certain tradeoffs in favoring certain combinations of traits over others (see also Dall et al., 2004; Bell, 2007). While our model was not developed initially within the context of behavioral syndromes, it can easily fit into this framework. Our work was initiated by questions related to the evolution of learning and social behaviors, and it is the first to address the evolution of different learning strategies, their possible co-evolution with the producer/scrounger social foraging strategies, and to show the emergence of a tradeoff favoring particular combinations of social and learning strategies.

Finally, it should be clear that the present analysis only outlines a basic framework for future work. It illustrates that persistent variation in cognitive abilities may be explained by their co-evolution with a mixture of social foraging strategies. However, it does not capture many additional processes that are likely to be involved in a realistic producer-scrunner game. First, there is evidence that individuals' tendencies to behave as producers or scrunchers can be shaped by experience (Katsnelson et al., 2008). This suggests an interesting, yet unexplored, interaction between learning rules for finding food and learning rules for choosing among strategies in a game. Second, if part-time scrunchers use social learning by paying attention to various cues related to food when they scrounge from others, many more learning steps may be available for complex learners. The integration of self and social learning may not be simple, however, since the reliability of social learning may be reduced by several different factors (see Laland, 2004 for a review). However, this initial analysis suggests that there is much to gain by studying the evolution of learning mechanisms within the social context under which they are operating.

Acknowledgments

The authors wish to thank E. Borenstein, L. Lehmann, E. Katsnelson, U. Grodzinski and two anonymous reviewers for their very helpful comments. This study was supported by the United States-Israel Binational Science Foundation Grant No. 2004412.

References

- Aoki K, Wakano JY, Feldman MW. The emergence of social learning in a temporally changing environment: A theoretical model. *Curr. Anthropol.* 2005; 46:334–340.
- Arbilly M, Motro U, Feldman MW, Lotem A. Recombination and the evolution of coordinated phenotypic expression in a frequency-dependent game. Submitted.
- Barnard CJ, Sibly RM. Producers and scroungers - a general model and its application to captive flocks of house sparrows. *Anim. Behav.* 1981; 29:543–550.
- Bateson M, Kacelnik A. Starlings' preferences for predictable and unpredictable delays to food. *Anim. Behav.* 1997; 53:1129–1142. [PubMed: 9236010]
- Beauchamp G. Learning rules for social foragers: Implications for the producer-scrounger game and ideal free distribution theory. *J. Theor. Biol.* 2000; 207:21–35. [PubMed: 11027477]
- Beauchamp G, Kacelnik A. Effects of the knowledge of partners on learning rates in zebra finches *Taeniopygia guttata*. *Anim. Behav.* 1991; 41:247–253.
- Bell AM. Future directions in behavioural syndromes research. *Proc. R. Soc. B.* 2007; 274:755–761.
- Ben-Akiva, M.; Lerman, S. Discrete Choice Analysis: Theory and Application to Travel Demand. Cambridge: MIT Press; 1985.
- Bereby-Meyer Y, Erev I. On learning to become a successful loser: A comparison of alternative abstractions of learning processes in the loss domain. *J. Math. Psychol.* 1998; 42:266–286. [PubMed: 9710551]
- Bergman A, Feldman MW. On the evolution of learning: Representation of a stochastic environment. *Theor. Popul. Biol.* 1995; 48:251–276.
- Bernstein C, Kacelnik A, Krebs JR. Individual decisions and the distribution of predators in a patchy environment. *J. Anim. Ecol.* 1988; 57:1007–1026.
- Borenstein E, Feldman MW, Aoki K. Evolution of learning in fluctuating environments: When selection favors both social and exploratory individual learning. *Evolution.* 2008; 62:586–602. [PubMed: 18182074]
- Bouchard TJ, McGue M. Familial studies of intelligence - a review. *Science.* 1981; 212:1055–1059. [PubMed: 7195071]
- Boyd, R.; Richerson, PJ. Culture and the Evolutionary Process. Chicago: University of Chicago Press; 1985.
- Boyd, R.; Richerson, PJ. An evolutionary model of social learning: The effect of spatial and temporal variation. In: Zentall, T.; Galef, BG., editors. *Social Learning*. Hillsdale: Erlbaum; 1988. p. 29-48.
- Boyd R, Richerson PJ. Why does culture increase human adaptability? *Ethol. Sociobiol.* 1995; 16:125–143.
- Burger JM, Kolss M, Pont J, Kawecki TJ. Learning ability and longevity: A symmetrical evolutionary trade-off in *Drosophila*. *Evolution.* 2008; 62:1294–1304. [PubMed: 18363867]
- Bussemeyer JR, Myung IJ. An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *J. Exp. Psychol. Gen.* 1992; 121:177–194.
- Caldwell CA, Whiten A. Scrounging facilitates social learning in common marmosets, *Callithrix jacchus*. *Anim. Behav.* 2003; 65:1085–1092.
- Camerer C, Ho TH. Experience-weighted attraction learning in normal form games. *Econometrica.* 1999; 67:827–874.
- Dall SRX, Houston AI, McNamara JM. The behavioural ecology of personality: Consistent individual differences from an adaptive perspective. *Ecol. Lett.* 2004; 7:734–739.
- Dernell J, March JG. Adaptation as information restriction: The hot stove effect. *Organ. Sci.* 2001; 12:523–538.

- Dukas, R. Constraints on information processing and their effects on behavior. In: Dukas, R., editor. *Cognitive Ecology: The Evolutionary Ecology of Information Processing and Decision Making*. Chicago: University of Chicago Press; 1998. p. 89-127.
- Dukas R. Costs of memory: Ideas and predictions. *J. Theor. Biol.* 1999; 197:41–50. [PubMed: 10036206]
- Dukas R, Duan JJ. Potential fitness consequences of associative learning in a parasitoid wasp. *Behav. Ecol.* 2000; 11:536–543.
- Eliassen S, Jorgensen C, Mangel M, Giske J. Quantifying the adaptive value of learning in foraging behavior. *Am. Nat.* 2009; 174:478–489. [PubMed: 19694535]
- Feldman MW, Aoki K, Kumm J. Individual versus social learning: Evolutionary analysis in a fluctuating environment. *Anthropol. Sci.* 1996; 104:209–231.
- Fritz J, Kotrschal K. Social learning in common ravens, *Corvus corax*. *Anim. Behav.* 1999; 57:785–793. [PubMed: 10202087]
- Galef BG, Laland KN. Social learning in animals: Empirical studies and theoretical models. *Bioscience.* 2005; 55:489–499.
- Giraldeau LA, Lefebvre L. Scrounging prevents cultural transmission of food-finding behavior in pigeons. *Anim. Behav.* 1987; 35:387–394.
- Giraldeau, LA.; Caraco, T. *Social Foraging Theory*. Princeton: Princeton University Press; 2000.
- Giraldeau LA, Dubois F. Social foraging and the study of exploitative behavior. *Adv. Stud. Behav.* 2008; 38:59–104.
- Groß R, Houston AI, Collins EJ, McNamara JM, Dechaume-Moncharmont FX, Franks NR. Simple learning rules to cope with changing environments. *J. R. Soc. Interface.* 2008; 5:1193–1202. [PubMed: 18337214]
- Hamblin S, Giraldeau L-A. Finding the evolutionarily stable learning rule for frequency-dependent foraging. *Anim. Behav.* 2009; 78:1343–1350.
- Harley CB. Learning the evolutionarily stable strategy. *J. Theor. Biol.* 1981; 89:611–633. [PubMed: 7289638]
- Houston AI, Sumida BH. Learning rules, matching and frequency-dependence. *J. Theor. Biol.* 1987; 126:289–308.
- Johnston TD. Selective costs and benefits in the evolution of learning. *Adv. Stud. Behav.* 1982; 12:65–106.
- Katsnelson E, Motro U, Feldman MW, Lotem A. Early experience affects producer-scrounger foraging tendencies in the house sparrow. *Anim. Behav.* 2008; 75:1465–1472.
- Kerr B, Feldman MW. Carving the cognitive niche: Optimal learning strategies in homogeneous and heterogeneous environments. *J. Theor. Biol.* 2003; 220:169–188. [PubMed: 12468290]
- Laland KN. Social learning strategies. *Learn. Behav.* 2004; 32:4–14. [PubMed: 15161136]
- Lefebvre L, Helder R. Scrounger numbers and the inhibition of social learning in pigeons. *Behav. Processes.* 1997; 40:201–207.
- Lendvai AZ, Barta Z, Liker A, Bokony V. The effect of energy reserves on social foraging: Hungry sparrows scrounge more. *Proc. R. Soc. B.* 2004; 271:2467–2472.
- Liker A, Barta Z. The effects of dominance on social foraging tactic use in house sparrows. *Behaviour.* 2002; 139:1061–1076.
- Liker A, Bokony V. Larger groups are more successful in innovative problem solving in house sparrows. *Proc. Natl. Acad. Sci. USA.* 2009; 106:7893–7898. [PubMed: 19416834]
- Lotem A. Learning to recognize nestlings is maladaptive for cuckoo *Cuculus canorus* hosts. *Nature.* 1993; 362:743–745.
- March JG. Learning to be risk averse. *Psychol. Rev.* 1996; 103:309–319.
- McNamara JM, Houston AI. Memory and the efficient use of information. *J. Theor. Biol.* 1987; 125:385–395. [PubMed: 3657218]
- Mery F, Kawecki TJ. Experimental evolution of learning ability in fruit flies. *Proc. Natl. Acad. Sci. USA.* 2002; 99:14274–14279. [PubMed: 12391295]
- Mery F, Kawecki TJ. A fitness cost of learning ability in *Drosophila melanogaster*. *Proc. R. Soc. B.* 2003; 270:2465–2469.

- Mery F, Kawecki TJ. An operating cost of learning in *Drosophila melanogaster*. *Anim. Behav.* 2004; 68:589–598.
- Mery F, Belay AT, So AKC, Sokolowski MB, Kawecki TJ. Natural polymorphism affecting learning and memory in *Drosophila*. *Proc. Natl. Acad. Sci. USA.* 2007; 104:13051–13055. [PubMed: 17640898]
- Niv Y, Joel D, Meilijson I, Ruppin E. Evolution of reinforcement learning in uncertain environments: A simple explanation for complex foraging behaviors. *Adapt. Behav.* 2002; 10:5–24.
- Plomin R. The genetics of *g* in human and mouse. *Nat. Rev. Neurosci.* 2001; 2:136–141. [PubMed: 11252993]
- Raine NE, Chittka L. The correlation of learning speed and natural foraging success in bumble-bees. *Proc. R. Soc. B.* 2008; 275:803–808.
- Rodriguez-Girones MA, Lotem A. How to detect a cuckoo egg: A signal-detection theory model for recognition and learning. *Am. Nat.* 1999; 153:633–648.
- Rogers AR. Does biology constrain culture? *Am. Anthropol.* 1988; 90:819–831.
- Shafir S, Reich T, Tsur E, Erev I, Lotem A. Perceptual accuracy and conflicting effects of certainty on risk-taking behaviour. *Nature.* 2008; 453:917–920. [PubMed: 18548069]
- Sibly, RM. Models of producer/scrounger relationships between and within species. In: Barnard, CJ., editor. *Producers and Scroungers: Strategies of exploitation and Parasitism*. Kent: Croom Helm; 1984. p. 267–287.
- Sih A, Bell AM. Insights for behavioral ecology from behavioral syndromes. *Adv. Stud. Behav.* 2008; 38:227–281.
- Sih A, Bell A, Johnson JC. Behavioral syndromes: An ecological and evolutionary overview. *Trends Ecol. Evol.* 2004a; 19:372–378. [PubMed: 16701288]
- Sih A, Bell AM, Johnson JC, Ziemba RE. Behavioral syndromes: An integrative overview. *Q. Rev. Biol.* 2004b; 79:241–277. [PubMed: 15529965]
- Stephens DW. On economically tracking a variable environment. *Theor. Popul. Biol.* 1987; 32:15–25.
- Stephens DW. Change, regularity, and value in the evolution of animal learning. *Behav. Ecol.* 1991; 2:77–89.
- Tolman EC. The inheritance of maze-learning ability in rats. *J. Comp. Psychol.* 1924; 4:1–18.
- Tracy ND, Seaman JW. Properties of evolutionarily stable learning rules. *J. Theor. Biol.* 1995; 177:193–198. [PubMed: 8558906]
- Tryon RC. Genetic differences in maze-learning ability in rats. *Yearbook of the National Society for Studies in Education.* 1940; 39:111–119.
- van Oers K, de Jong G, van Noordwijk AJ, Kempenaers B, Drent PJ. Contribution of genetics to the study of animal personalities: A review of case studies. *Behaviour.* 2005; 142:1185–1206.
- Wakano JY, Aoki K, Feldman MW. Evolution of social learning: A mathematical analysis. *Theor. Popul. Biol.* 2004; 66:249–258. [PubMed: 15465125]
- Wolf M, van Doorn GS, Weissing FJ. Evolutionary emergence of responsive and unresponsive personalities. *Proc. Natl. Acad. Sci. USA.* 2008; 105:15825–15830. [PubMed: 18838685]
- Wolf M, van Doorn GS, Leimar O, Weissing FJ. Life-history trade-offs favour the evolution of animal personalities. *Nature.* 2007; 447:581–584. [PubMed: 17538618]

(a)

(b)

Fig. 1.

Success of the three learning rules in relation to the number of steps in a game. Each line represents the relative weight (in %) of each of the learning alleles among agents whose payoff was in the top 50% after J steps, averaged across 100,000 simulated populations. (a) F_1 populations (pure producers). (b) F_2 populations (part-time scroungers). Memory factor $\alpha = 0.5$.

Fig. 2.

Mean haplotype frequency for different costs to complex learning (γ), and for different number of steps in a game (J). Black columns represent haplotypes associating the learning alleles L_1 (complex), L_2 (simple) and L_0 (no learning) with the producer allele F_1 , and grey bars represent their association with the part-time scrounger allele F_2 . Mean for each haplotype is calculated for generations 1500–3000, for 10 repeats of the simulation. Error bars represent the standard error.

Fig. 3. Example of evolutionary dynamics in a population with generation lifespan $J = 50$ steps. Population size is $n = 300$, mutation rate in both genetic loci is $\mu = 1/n$. The frequency of each haplotype at each generation is indicated by the area its color occupies on the Y axis. (a) Complex learning cost $\gamma = 0$. (b) Complex learning cost $\gamma = 0.01$.

Fig. 4. ΔF_1 and ΔF_2 as a function of the complex learning cost, γ . For each value of γ , the corresponding values of ΔF_i are averaged over all J (the number of steps). Note that for γ in the vicinity of 0.01, $\Delta F_1 > 0$, while $\Delta F_2 < 0$.

Food patches parameters and expected value. E_{12} and E_{34} represent perceptual unification of two patches based on a common visual cue (e.g. color or shape); the black square, black triangle, white circle and white pentagon demonstrate possible visual characteristics of the patches which may cause perceptual unification based on color. The shapes are shown here for illustrative purposes only.

Table 1

patch	E_1 ■	E_2 ▲	E_3 ○	E_4 ⬡	E_{12} ■▲	E_{34} ○⬡
parameter						
Nutritional value	4	1.5	0.75	0.25		
Probability of finding food	0.25	0.33	0.5	1		
Expected value	1	0.5	0.375	0.25	0.75	0.313

Table 2

Symbols used in the simulations and their descriptions

Symbol	Description
F_1	Foraging allele: pure producer (1.0 probability to apply producer strategy)
F_2	Foraging allele: part-time scrounger (0.5 probability to apply producer strategy)
L_0	Learning allele: non-learning
L_1	Learning allele: complex learning rule
L_2	Learning allele: simple learning rule
n	Number of agents in the population
J	Number of steps in the game (equivalent to agents' lifespan)
G	Number of generations
μ	Mutation rate
γ	Cost of complex learning rule (% deducted from final cumulative payoff)
α	memory factor (weight of past events in memory when updating a new experience)