

# Social learning and information sharing: An evolutionary simulation model of foraging in Norway rats

Jason Noble, Elio Tuci and Peter M. Todd  
Center for Adaptive Behavior and Cognition  
Max Planck Institute for Human Development  
Lentzeallee 94, D-14195 Berlin, Germany

## Abstract

Social learning is distinguished from innate behaviour and individual learning as a behavioural strategy. We investigate simple mechanisms for social learning in an evolutionary simulation of food-preference copying in Norway rats. These animals learn preferences by interacting with conspecifics, but, unexpectedly, they fail to learn aversions after interacting with a poisoned demonstrator. They also follow each other to food sites. Simulation results show that failure to discriminate between sick and healthy demonstrators may be due to details of food toxicity in foraging environments. A seemingly complex instance of social information transmission is explained through the action of simple behaviours in an appropriately structured environment.

## 1 Introduction

Animal behaviour can be seen as the problem of what to do next, or as Maes (1989) put it, “how to do the right thing”. Natural selection is the force that has shaped the behavioural strategies of the animals we see today, but, clearly, natural selection has arrived at different solutions in different species. For an animal facing a particular environmental challenge, three broad sources of strategy can be distinguished: instinct, learning, and social learning. For example, suppose that a foraging animal has to decide whether or not to eat a piece of toxic, unripe fruit that it has found. The decision might be made instinctively: the animal has an inherited tendency to avoid fruit of that colour and texture. Alternatively, the animal might have learned through bitter experience that such fruit is unpalatable. Finally, it could have learned socially:

perhaps it has observed conspecifics rejecting this kind of fruit, or has seen conspecifics become ill after eating it. In this paper we focus on the third strategy source, social learning, to explore the way in which simple specific mechanisms of social information gathering can interact with structured environments to yield unexpected behavioural implications.

The idea that animals sometimes learn from the behaviour of others poses both why and how questions, i.e., questions of evolved function and questions of mechanism. In recent years there has been some progress towards understanding the adaptive function of social learning. Models of cultural transmission (Boyd & Richerson, 1985), highly horizontal transmission (Laland, Richerson, & Boyd, 1996), and what economists call herding behaviour (Bikhchandani, Hirshleifer, & Welch, 1998) help to delineate the conditions under which it will be advantageous for individuals to learn from others rather than finding things out for themselves. However, these models tend to result in rather general conclusions. For example, Laland et al. (1996, p. 140) summarize the results of work on cultural transmission thus:

When environments change very slowly, all information should be transmitted genetically, since the modest demands for updating are easily met by the genetic system responding to selection. When environmental change is very rapid, tracking by pure individual learning is favored. At intermediate rates of change, social learning is an advantage.

Results like these are useful, but to get a complete picture of any one instance of social learning we also need to understand the mechanism involved: *how* exactly do the animals gain information from the behaviour of their conspecifics? A variety of mechanisms have been proposed (see Galef,

1988, for a review). Imitation is one type of social learning that has received a great deal of attention: the idea is that certain cognitively sophisticated animals might “from an act witnessed learn to do an act” (Thorndike, 1911). However, the precise way in which this feat could be achieved has never been properly spelled out. As is usually the case, a seemingly complex behaviour may be produced by a range of simpler mechanisms. For example, Galef (1988) discusses “stimulus enhancement” in which a tendency on the part of naive animals to approach conspecifics leads to their being more likely to encounter one set of stimuli rather than another, and thus shapes their (individual) learning experience. Galef also discusses “contagious behaviour”, in which the performance of one behaviour, e.g., feeding, acts as a trigger for the performance of the same behaviour by others. No long-term learning is necessarily involved in contagious behaviour, but it is nevertheless a means by which animals could gain information from their conspecifics.

The power of these simpler mechanisms has probably been underestimated. Several phenomena that were once seen as involving imitation, such as the opening of milk bottles by birds (Fisher & Hinde, 1949) and the washing of food by monkeys (Kawamura, 1959), have since been questioned (Sherry & Galef, 1984; Galef, 1988; Cheney & Seyfarth, 1990; Byrne, 1995), and parsimonious explanations have been offered in terms of processes like stimulus enhancement. However, it is difficult to design experiments that conclusively expose the mechanism at work in particular cases. Especially with regard to nonhuman primates, controversy continues over just how sophisticated animal social learning is (Byrne, 1995; Tomasello, 1996; Heyes, 1998).

We believe that the individual-based simulations characteristic of artificial life can be useful tools in the investigation of social learning, much as they have been for studying the evolution of individual learning (e.g., Todd & Miller, 1991; Belew & Mitchell, 1996). It has long been recognized within artificial life that complex global phenomena can arise from simple local rules, and this is precisely what some researchers suspect is happening in animal social learning: individuals follow a simple rule (e.g., “stay close to your mother”) and, in combination with associative learning, the overall pattern of behaviour that arises makes human observers suspect imitation. Although work in artificial life has certainly considered social dynamics in contexts such as foraging (Bonabeau & Cogne, 1996), communication (Di Paolo, 1997), and flocking or schooling (Reynolds, 1987; Terzopoulos, Tu,

& Grzeszczuk, 1995), there has been relatively little work on the specific topic of social learning. The model most relevant to our own work is by Toquenaga, Kajitani, and Hoshino (1995), who constructed a simulation of foraging behaviour in egrets. Toquenaga et al. demonstrate that stimulus enhancement is an important determinant of the evolution of flock foraging and colonial roosting, and is more likely to evolve when resources are patchy rather than evenly distributed.

We look at social learning in Norway rats (*Rattus norvegicus*)—an opportunistic, central place foraging species—to see how their specific and rather surprising social learning mechanisms may have evolved in response to environmental features. These rats employ at least two simple mechanisms that allow them to share information about food (Lore & Flannelly, 1977; Galef, 1996). Firstly, they have a robust tendency to copy the feeding preferences of their conspecifics. A rat will develop a marked preference for a novel food that it smells on the breath of another, and the effect is strong enough to make a rat choose the novel food type over its normal diet, despite the fact that rats usually avoid new foods. Note that this is not imitation, but the application of a simple behavioural rule: “if others have been eating  $X$ , then  $X$  is good”. The acquired preference is specific to the context of food; the rat will have no special preference for nesting sites or nesting materials that carry the same odor. The key stimulus is the detection of the novel food odor in combination with carbon disulfide, a component of rat breath. Rats will not, for instance, develop a preference for foods that an experimenter has wiped onto the fur of another rat. Secondly, rats will spontaneously follow conspecifics on foraging trips out of the nest; this habit is especially pronounced in younger animals. Such behaviour clearly suggests that stimulus enhancement may be occurring. One of the ways a rat could come to exploit a new food source would be simply by following another and learning from the experience.

Galef, Wigmore, and Kennett (1983) uncovered an apparent paradox in rats’ social learning. They assumed that if rats could acquire food preferences through social interaction, they would probably also be able to learn an aversion to a novel, toxic food by smelling it on the breath of a demonstrator rat and simultaneously noting that the demonstrator was suffering from acute gastro-intestinal distress—that is, it was sick. Experiments revealed, however, that this is not the case. Rats are not sensitive to a demonstrator’s state of health, and

in fact only ever develop a *preference* for the novel food.

This surprising finding was the starting point for our own investigations. The adaptive value of distributed intelligence through copying conspecifics' food preferences seems clear in an opportunistic forager that must deal with new and potentially toxic foods, especially when seen alongside the fact that rats will normally avoid novel foods. But why don't they discriminate between sick and healthy demonstrators? It is not because they can't—rats perform well on a wide variety of discrimination tasks, and are capable of identifying sick conspecifics using odor and behavioural cues (Lavin, Freise, & Coombes, 1980; Bond, 1982). Curiously, other species such as blackbirds (Mason, Arzt, & Reiding, 1984) and chickens (Johnston, Burne, & Rose, 1998) *do* manage to learn both preferences and aversions through observation.

We suspected that the answer might depend on characteristics of the rats' foraging environment; specifically on the probability that eating a toxic food would result in the death of the animal. To test this suspicion we constructed an evolutionary simulation, within which we systematically varied the lethality of toxic foods in the environment, and observed the effect on the evolution of a gene for discriminating between sick and healthy demonstrators. We were also interested in following behaviour, and extended the initial model to include this possibility. In particular, we wanted to demonstrate that these two simple mechanisms—copying and following—could together account for apparently complex social learning. We were also interested in determining whether there was any interaction between the two: are the benefits of copying food preferences and those of following others independent, or do they interfere with each other in some way?

## 2 The evolution of preference copying

### 2.1 Modelling learning rats

In our simulation an initial population of 100 rats foraged from five centrally located nests. Rats forage by night, and we divided each night into five foraging periods. During each period, a rat could visit one of 25 foraging sites: if it found food at that site, it had to decide whether or not to eat it. If a rat chose to eat, and if there was sufficient food, it would fill its stomach immediately and return

to the nest—the night's foraging was over. However, rats that rejected the food they had found, and rats that ate only a partial meal due to competition, would continue to forage until they had eaten their fill or until all five periods had passed. If a rat consumed nutritious food it gained energy, but 10% of the food types in the environment were toxic. If a rat consumed a toxic food there was a parameter governing the probability that it would die at once; otherwise it would lose some energy and would show signs of poisoning when it returned to the nest. Rats that consumed no food at all would eventually die of starvation.

The rats were given a simple memory: any food they encountered would either be novel, familiar or aversive. For newborn rats, all foods were novel. The first variable under evolutionary control was the probability  $E$  that a rat would eat a novel food. After eating a new food, the rat would remember it as either familiar or aversive depending on whether or not it was toxic. Rats would always eat familiar foods, and always reject aversive foods. Rats were also assumed to be capable of remembering where they had foraged last night, and a binary gene  $P$  controlled whether or not they would persist and return to their last successful feeding site, if any, during the first period of the subsequent evening's foraging.

Each individual had two more genes controlling its preference-copying behaviour: a binary gene  $S$  indicating whether or not the rat would learn about new foods by smelling the breath of other rats, and a binary gene  $D$  that controlled whether or not the artificial rat (unlike real rats) would discriminate with regard to the state of health of another rat whose breath it smelled, avoiding foods eaten by sick rats and preferring foods on the breath of healthy rats. Thus, upon returning to the nest after each foraging trip, rats could potentially smell the breath of a randomly chosen nestmate as a way of gaining information about new foods in the environment. If a rat had only the smell-based learning ability,  $S$ , it would simply smell the breath of a nestmate and become familiar with the type of food the nestmate had eaten that day, if any. But the rat might thereby develop familiarity with and hence preference for a food that was in fact toxic: some rats died immediately after eating a poisonous food, but others made it back to the nest and were ill. A rat that learned such a preference for toxic food from a conspecific that survived the poisoning might not itself be so lucky when it consumed that food. Only if the learning rat also had the ability to discriminate,  $D$ , would it develop prefer-

ence or aversion for the new food type depending on whether the nestmate was showing signs of poisoning. (A low level of error could also be associated with this discrimination ability.)

When a rat had accumulated a certain amount of energy, it would undergo simplified asexual reproduction. The carrying capacity of the environment was fixed at 100. This implied that if a newborn individual could not take the place of a rat that had recently died of poisoning or starvation, the current oldest rat would be selected for death by old age in order to make room. Newborn rats inherited (with a small chance of mutation) the four-element behavioural strategy of their parent described above, plus a level of starting energy to ensure that they did not immediately starve.

Parcels of food appeared in random sites in the environment at a constant rate. In order to ensure that the rats always had to deal with novelty, a total of 100 food types were used, but with a “window” of 10 food types that could appear at any one time. Every 16 days the window would advance, so that one old food type stopped appearing, and a novel one entered the scene. The lifespan of an individual rat (which was an emergent property of the simulation) never grew long enough for it to experience all 100 food types.

While we have tried to make this model reflect part of the lives of real rats, the parameters used in the simulation<sup>1</sup> are not as closely matched to real data as we would like. There is a great deal of information available on the behaviour of rats in the laboratory, but data on the ecology of wild rats is not extensive (see e.g., Lore & Flannelly, 1977; Lore & Schultz, 1989). We have therefore pitched the simulation at a relatively abstract level, and in so doing we hope to have captured some aspects of the selection pressures impinging on social foragers in general. However, we also anticipate that the results of modelling will help us to reverse-engineer the environment of real rats, by demonstrating necessary connections between variables for which data is available and those for which it is not.

## 2.2 Results

In accordance with our hypothesis, we varied the probability of death due to eating a toxic food,

<sup>1</sup>Rat stomach capacity = 10 food units; cost of living = 1 unit per day; default error level in discrimination = 0.01; mutation rate = 0.05, standard deviation for mutation of real-numbered genes = 0.1; food parcel mean size = 100 units, standard deviation = 40, food parcel input rate = 5 per day; reproduction level = 1000 units, cost of offspring = 500 units, offspring starting energy = 400 units.

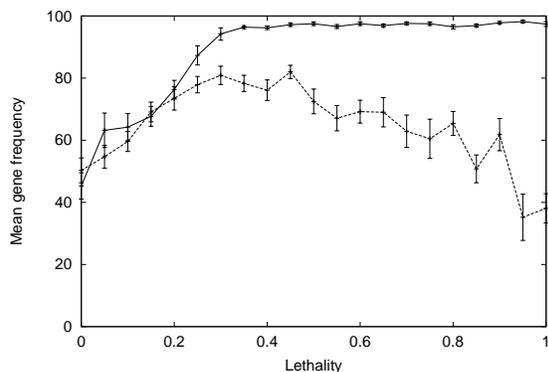


Figure 1: Mean frequency of the smelling and discrimination genes *S* and *D* by lethality level—solid and dashed lines respectively. In this and subsequent figures, error bars show the standard error across ten simulation runs.

to find out what difference this made to the behavioural strategies of the rats. For each level of lethality investigated, ten evolutionary runs, each 200,000 days in length, were performed. If some populations went extinct before this time limit was reached, as sometimes happened in the more lethal environments, the simulation was repeated until ten complete runs had been recorded. The statistics reported below describe the state of populations at the end of these runs.

When the lethality value was zero, i.e., in a benign environment, there was no selection pressure on the rats to smell each other or to discriminate: Figure 1 shows that gene frequencies for *S* and *D* remained close to 50%, the value expected by chance. Nor were the rats particularly persistent, as we will see later in Figure 7. At the same time, the mean probability for eating novel foods (*E*) was high at 89.8% (Figure 2). The mean lifetime of a rat was 268 days, during which time approximately 17 new foods would have made an appearance. The rats were, on average, familiar with 13.3 food types, and had aversions to 2.1 foods: these frequencies are not too far from the 10% base rate of toxic foods in the environment. Thus, when poison results only in a stomach ache and an aversion to trying that food again in future, simulated rats are open to trying new foods, and pay no special attention to the eating habits of others.

It is clear from Figure 1 that as the lethality level increases there is increasing selection pressure for learning from others. Rats become more likely to smell the breath of conspecifics, and to discriminate between the sick and the healthy when doing

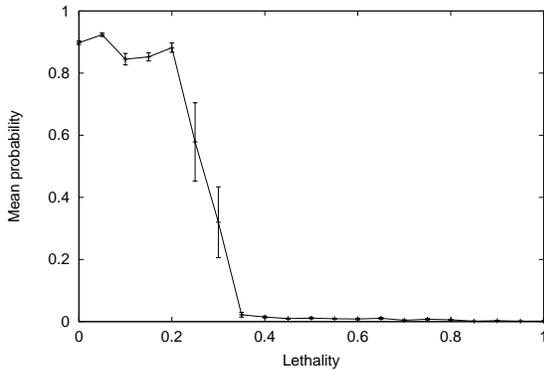


Figure 2: Mean value of the real-valued gene  $E$ , governing the probability of eating novel food, by lethality level.

so. At the same time, they are still willing to try new foods that they come across (Figure 2). However, when the lethality level reaches approximately 0.3, there is something like a phase transition in the results. The probability of eating novel foods drops dramatically; the overall mean probability for eating a novel food, across all runs with lethality greater than 0.3, was less than 1%. As lethality approaches 1.0, there is a uniform strong selection pressure on the gene for smelling the breath of others. Importantly, though, the selection pressure on discriminating decreases: as poison became more and more dangerous, it was no longer important to pay attention to the health of a demonstrator—we will argue in section 4 that this result explains the failure of real rats to discriminate. Aggregating the results for lethality levels greater than 0.3, the rats were familiar with about the same number of foods (10.0), but they developed almost no aversions (0.2). There were other effects: because the rats were more cautious in their treatment of novel foods, they generally ate less, which led to increases in the inter-birth interval and the average lifespan.

The finding that selection for discrimination falls off with increasing lethality could have something to do with the error rate of the rats' discriminative ability, normally set to 1%. After all, when lethality is equal to one, all rats that have been poisoned will die before returning to the nest. Under these circumstances there is not much point in discriminating, and errors in discrimination could be positively harmful. We therefore considered the evolution of the gene  $D$  under zero-error and 5% error conditions. Figure 3 shows that a higher error rate indeed leads to selection *against* discriminat-

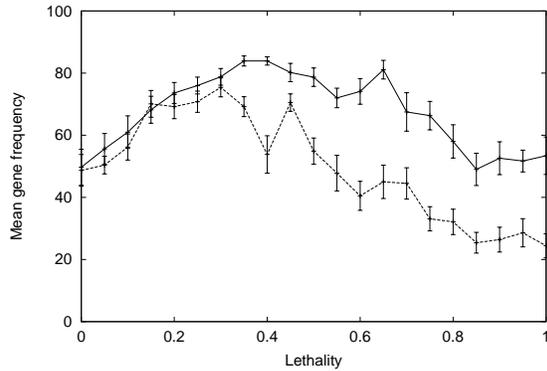


Figure 3: Mean frequency of the discrimination gene  $D$  given zero (solid line) and 5% (dashed line) error rates in discriminative ability.

ing. But even with zero error, the frequency of the gene  $D$  is close to the expected value for random drift (i.e., 50%), given high levels of lethality. Even if one's discrimination ability were perfect, it appears that it would not be a selective advantage in an environment where eating a poisonous food would kill you four times out of five.

### 3 Adding following to the model

#### 3.1 Changes in the model

In order to make following behaviour possible, the rats were given a fifth gene,  $F$ , that determined the probability that they would follow a random conspecific when leaving the nest on the first foraging period of the evening. Following was only ever selected if the rat was not being persistent. In other words, the decision tree for the rats was as follows.

- Check persistence gene,  $P$ ; if set, was I successful last night? If so, return to that site.
- If either condition was not met, check percentage chance to follow,  $F$ ; should I follow someone else?
- If neither of the above selected, choose a random site.

Rats that followed simply went to the same site as a randomly chosen conspecific who was not also following, but there was no way for a rat to tell whether it was following a persistent conspecific or one that had merely selected a site at random. In

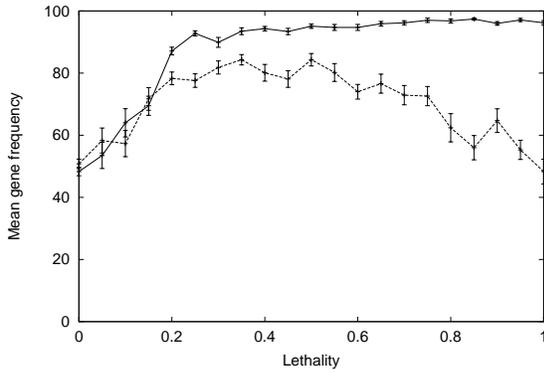


Figure 4: Mean frequency of the smelling and discrimination genes  $S$  and  $D$  in the model with following.

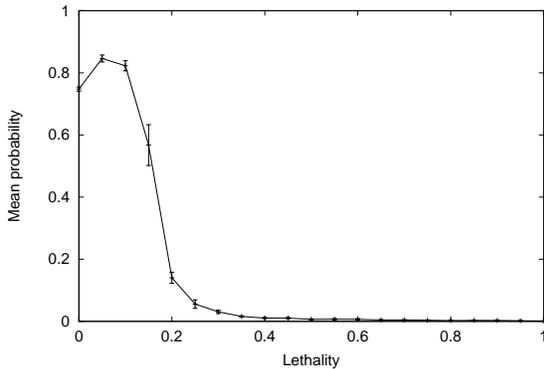


Figure 5: Mean value of the real-valued gene  $E$  governing the probability of eating novel food in the following case.

the unlikely event that *all* the rats chose to follow, the dubious possibility that they would all stay in the nest waiting for someone to leave was dealt with by sending them all to random sites. The extended model was in other respects the same as the basic model.

### 3.2 Results

In Figures 4 and 5 we can see that the evolution of the genes for smelling, discrimination and eating novel foods are not much affected by the possibility that rats may follow each other to food sites, although the transition between liberal and conservative attitudes to new foods occurs at a lower level of lethality.

The results for the evolution of the following gene are shown in Figure 6. We were somewhat surprised

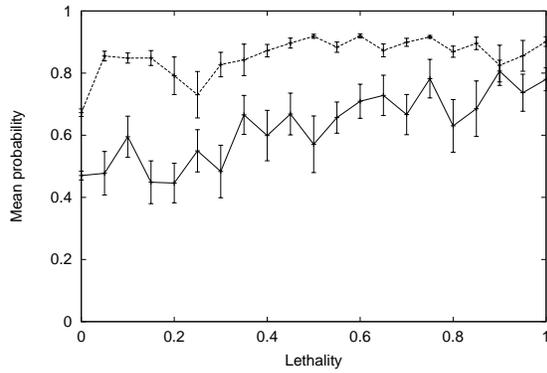


Figure 6: Mean value of the real-valued gene  $F$ , governing the probability of following other rats to a foraging site. The solid line shows the standard condition, while the dashed line shows the effect of a less uniform food distribution (see text).

to find that following other rats was not strongly selected for (solid line). There is a modest positive relationship between the mean value of  $F$  and lethality, but for the more benign environments the data could just as well be the result of a random walk. The large standard error values indicate that there was a great deal of variation in the final population mean value of  $F$  across runs.

To find circumstances that could select for following behaviour, we ran a variant of the simulation in which food was less uniformly distributed in the environment. Following others to food is a kind of stimulus enhancement, and Toquenaga et al. (1995) found that stimulus enhancement was more likely given patchy food distributions. Instead of five food parcels with a mean size of 100 units arriving per day, one food parcel with a mean size of 500 units was supplied; Figure 6 shows that under these conditions following behaviour was indeed more strongly selected for (dashed line).

Another way of looking at following behaviour is to consider its effects on the gene for persistence. Figure 7 shows the mean frequencies for the gene  $P$  for the initial model and for the model with following added. In the initial case (solid line), we can see that persistence becomes more important with increasing levels of lethality. In a dangerous world where you don't want to try anything new, it pays to go back to the site of yesterday's successful feeding. However, when following behaviour is possible (dashed line), there is not much selection pressure on persistence.

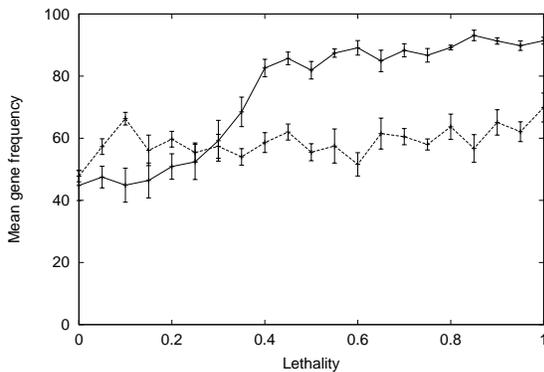


Figure 7: Comparison of the frequencies for the persistence gene  $P$  with (dashed line) and without (solid line) the possibility of following behaviour.

## 4 Implications for real rats

The failure of real rats to discriminate between sick and healthy demonstrators may be due to the details of food toxicity in their foraging environments. In the initial model, when the lethality value crosses the threshold of about 0.3 the whole pattern of rat behaviour changes. Rats become extremely wary of new foods, but attend closely to what their conspecifics are eating. However, because they “know” that their nestmates are just as conservative as they are about trying new and potentially dangerous foods, the need to discriminate between sick and healthy demonstrators is reduced. Observing a poisoned demonstrator would doubtless provide useful information, but it becomes such a rare event that there is little selection pressure for paying attention to it.

The possibility of error in discrimination just makes things worse. In the terminology of signal detection theory, a hit would be correctly identifying a food as poisonous after observing a sick demonstrator who has eaten that food. A miss would be failing to identify a food as poisonous under these circumstances. Misses would certainly be costly, but false alarms may at times be even more so: given that one has a low likelihood of eating new foods, and that there is a limited number of foods present in the environment, believing falsely that a palatable food is poisonous could deprive a rat of a much-needed food source. As the error rate for discrimination increases, there must come a point when it is better to simply accept all foods detected on the breath of others, thus risking occasional poisoning but ensuring that no palatable food source ever goes unexploited. It is not clear

what the error rates for discrimination would be in real rats, but if we recognize that animals can be ill for reasons other than food poisoning, and that they may conceal illness when it exists, the levels of 1% and 5% seem conservative.

Furthermore, discrimination in the real world would not come for free. The animal would have to pay the time and energy costs of developing a sensory system and decision mechanism that allowed it to detect sickness in others and then act accordingly; the blind acceptance of any food odor that is smelled in conjunction with carbon disulfide is clearly a simpler way to do things. In the simulation we did not attempt to model this fact, i.e., we did not include any direct costs on the ability to discriminate. The existence of such direct costs for real rats would make discrimination even less likely to evolve.

There is a clear empirical prediction arising from the initial model. Firstly, if we consider foraging populations with very low levels of danger associated with eating toxic food, animals in these populations should be content to eat novel food and ignore the experiences of others. Next, as lethality increases, individuals will become more likely to pay attention to the eating habits of others, and to discriminate as to the state of health of demonstrators. Finally, above a certain threshold of lethality we expect to find animals with a very low likelihood of eating novel foods, a great interest in what others are eating, and no strong tendency to discriminate between sick and healthy demonstrators. Norway rats appear to fit this last profile. Blackbirds (Mason et al., 1984) and chickens (Johnston et al., 1998) may fit the second, while specialist species probably fit the first.

The results for the extended simulation with the possibility of following have less clear implications. In increasingly lethal environments, following others to food sources is moderately favoured. When the distribution of food is patchy, there is reasonably strong selection pressure for following behaviour; this is in accordance with Toquenaga et al.’s (1995) findings that patchy food distributions will promote stimulus enhancement behaviours. But the interaction between following and persistence (see Figure 7) demands further analysis. The problem is that mean values for the gene that governs the probability of following in the non-patchy environment sit close to 50% and have large standard error values over simulation runs; this is exactly what would be expected if there was no selection pressure on the gene whatsoever, and suggests the conclusion that following is not partic-

ularly adaptive under these circumstances. Nevertheless, the observed means imply that about half the time rats will be following a conspecific rather than choosing a random food site, assuming they are not being persistent and returning to the scene of yesterday's success. Comparing the initial model to the model with following added, we find that the gene for persistence is selected for in the former case, but not in the latter. The apparently chance levels of following behaviour are affecting the evolution of persistence; one could even argue that following behaviour *takes the place* of persistence. It is therefore not so clear that following behaviour is adaptively neutral.

Presumably there is a trade-off between following others to food but having to share it when you get there, and thus taking a chance that you will get less than a full meal, versus choosing a food site randomly and having it to yourself. (However, there were always more rats than food sites; on average a rat would be sharing with four others in any case.) This may also have something to do with the results for persistence: if some fraction of the population is going to follow you to your current favourite food site, the benefits of persistence may be swamped by the costs of attracting a large and hungry crowd. Clarification of these issues must wait for future work.

## 5 Conclusions

We see foraging behaviour in rats as a paradigm case in which a seemingly complex instance of social (or even cultural) information transmission can be explained through the action of simple behaviours in an appropriately structured environment—an instance of ecological rationality (Gigerenzer, Todd, & The ABC Research Group, 1999). In fact, the behaviour underlying rat food preference copying is even simpler than was originally expected by those looking for simple mechanisms. That is, rats pay no attention to whether the individual they are learning a preference from is suffering from food poisoning. Our model has added to this picture by showing that this strategy may well have evolved just because eating poison lessens the likelihood that a rat will survive to influence any of its conspecifics. Similarly, when following behaviour is possible and occurs at arbitrary levels in the population, rats do not even need to remember the site where they found food the night before.

We plan to investigate further how environment structure could have affected the evolution of different forms of learning in rats. As indicated in

section 1, Laland et al. (1996) have made general predictions as to when social and individual learning will evolve with respect to the rate of environmental change. We will extend our more specific individual-based model to see whether this prediction holds in the case of rat food-preference learning.

A particularly interesting extension to this work would be to add possibilities for selfishness and deception on the part of individual rats. It is conceivable, for example, that rats might evolve the ability to feign sickness in order to convince others that a new food source was unpalatable, and thus keep it for themselves. More simply, they might avoid contact with other rats in the nest so that no-one else came to know what they had been eating. Much of the recent artificial life work on communication (see e.g., Bullock, 1997; Noble, 1998) addresses this sort of question, i.e., whether and how honest communication can be maintained in the face of possible conflicts of interest. Although it is not normally considered as such, social transfer of information about food can clearly be a kind of communication; we hope to integrate the two perspectives in subsequent models.

## Acknowledgements

Thanks to Seth Bullock and Jennifer N. Davis for useful discussions. E.T. would like to thank the Italian government for financial support.

## References

- Belew, R., & Mitchell, M. (Eds.). (1996). *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Addison-Wesley, Reading, MA.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1998). Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives*, 12(3), 151–170.
- Bonabeau, E., & Cogne, F. (1996). Oscillation-enhanced adaptability in the vicinity of a bifurcation: The example of foraging in ants. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pp. 537–544. MIT Press / Bradford Books, Cambridge, MA.
- Bond, N. W. (1982). Transferred odor aversions in adult rats. *Behavioral and Neural Biology*, 35, 417–421.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. The University of Chicago Press, Chicago.
- Bullock, S. (1997). An exploration of signalling behaviour by both analytic and simulation means for both discrete and continuous models. In Husbands & Harvey (1997), pp. 454–463.

- Byrne, R. W. (1995). *The Thinking Ape: Evolutionary Origins of Intelligence*. Oxford University Press, Oxford.
- Cheney, D. L., & Seyfarth, R. M. (1990). *How Monkeys See the World*. University of Chicago Press, Chicago.
- Di Paolo, E. A. (1997). Social coordination and spatial organization: Steps towards the evolution of communication. In Husbands & Harvey (1997), pp. 464–473.
- Fisher, J., & Hinde, R. A. (1949). The opening of milk bottles by birds. *British Birds*, *42*, 347–357.
- Galef, Jr., B. G. (1988). Imitation in animals: History, definition, and interpretation of data from the psychological laboratory. In Zentall, T. R., & Galef, Jr., B. G. (Eds.), *Social Learning: Psychological and Biological Perspectives*, pp. 3–28. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Galef, Jr., B. G. (1996). Social enhancement of food preferences in Norway rats: A brief review. In Heyes & Galef (1996), pp. 49–64.
- Galef, Jr., B. G., Wigmore, S. W., & Kennett, D. J. (1983). A failure to find socially mediated taste aversion learning in Norway rats (*R. norvegicus*). *Journal of Comparative Psychology*, *97*, 358–363.
- Gigerenzer, G., Todd, P. M., & The ABC Research Group (1999). *Simple Heuristics That Make Us Smart*. Oxford University Press, New York.
- Heyes, C. M. (1998). Theory of mind in nonhuman primates. *Behavioral and Brain Sciences*, *21*, 101–148.
- Heyes, C. M., & Galef, Jr., B. G. (Eds.). (1996). *Social Learning in Animals: The Roots of Culture*. Academic Press, San Diego, CA.
- Husbands, P., & Harvey, I. (Eds.). (1997). *Proceedings of the Fourth European Conference on Artificial Life (ECAL'97)*. MIT Press / Bradford Books, Cambridge, MA.
- Johnston, A. N. B., Burne, T. H. J., & Rose, S. P. R. (1998). Observation learning in day-old chicks using a one-trial passive avoidance learning paradigm. *Animal Behaviour*, *56*, 1347–1353.
- Kawamura, S. (1959). The process of sub-culture propagation among Japanese macaques. *Primates*, *2*, 43–60.
- Laland, K. N., Richerson, P. J., & Boyd, R. (1996). Developing a theory of animal social learning. In Heyes & Galef (1996), pp. 129–154.
- Lavin, M. J., Freise, B., & Coombes, S. (1980). Transferred flavor aversions in adult rats. *Behavioral and Neural Biology*, *28*, 25–33.
- Lore, R., & Flannelly, K. (1977). Rat societies. *Scientific American*, *236*(5), 106–116.
- Lore, R. K., & Schultz, L. A. (1989). The ecology of wild rats: Applications in the laboratory. In Blanchard, R. J., Brain, P. F., Blanchard, D. C., & Parmigiani, S. (Eds.), *Ethoexperimental Approaches to the Study of Behavior*, Vol. 48 of *NATO Advanced Science Institutes series, Series D: Behavioral and Social Sciences*, pp. 607–622. Kluwer, Boston.
- Maes, P. (1989). How to do the right thing. *Connection Science*, *1*(3), 291–323.
- Mason, J. R., Arzt, A. H., & Reidinger, R. F. (1984). Comparative assessment of food preferences and aversions acquired by blackbirds via observational learning. *Auk*, *101*, 796–803.
- Noble, J. (1998). Evolved signals: Expensive hype vs. conspiratorial whispers. In Adami, C., Belew, R., Kitano, H., & Taylor, C. (Eds.), *Artificial Life VI*, pp. 358–367. MIT Press, Cambridge, MA.
- Reynolds, C. W. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, *21*(4), 25–34. SIGGRAPH '87 Conference Proceedings.
- Sherry, D. F., & Galef, Jr., B. G. (1984). Cultural transmission without imitation: Milk bottle opening by birds. *Animal Behaviour*, *32*, 937–938.
- Terzopoulos, D., Tu, X., & Grzeszczuk, R. (1995). Artificial fishes. *Artificial Life*, *1*(4), 327–351.
- Thorndike, E. L. (1911). *Animal Intelligence*. MacMillan, New York.
- Todd, P. M., & Miller, G. F. (1991). Exploring adaptive agency II: Simulating the evolution of associative learning. In Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pp. 306–315. MIT Press / Bradford Books, Cambridge, MA.
- Tomasello, M. (1996). Do apes ape?. In Heyes & Galef (1996), pp. 319–346.
- Toquenaga, Y., Kajitani, I., & Hoshino, T. (1995). Egrets of a feather flock together. *Artificial Life*, *1*(4), 391–411.