THE ROYAL SOCIETY

doi 10.1098/rspb.2000.1342

Explaining social learning of food preferences without aversions: an evolutionary simulation model of Norway rats

Jason Noble*, Peter M. Todd and Elio Tuci†

Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195, Berlin, Germany

Norway rats (*Rattus norvegicus*) transmit preferences for novel foods socially by smelling each other's breath. However, rats fail to learn aversions, acquiring a preference even if the rat whose breath they smell has been poisoned. Rats can distinguish between sick and healthy conspecifics and social learning of both preferences and aversions is present in other species—hence it is unclear why rats cannot learn aversions socially. We constructed an evolutionary simulation in which a population of rats foraged from a central location, exploiting food sites that could contain edible or toxic foodstuffs. We examined the relationship between toxin lethality and selection for individual versus social learning and discrimination between sick and healthy conspecifics in order to allow learning of both preferences and aversions. At low lethality levels individual learning was selected for and at intermediate levels we found social learning of both preferences and aversions. Finally, given high lethality levels the simulated rats would employ social learning but failed to learn aversions, matching the behaviour of real rats. We argue that Norway rats do not learn aversions socially because their environment may contain only highly lethal toxins which make interaction with a sick conspecific an extremely rare event.

Keywords: social learning; foraging; Norway rats; evolutionary simulation

1. INTRODUCTION

Opportunistic generalist foragers face a particular challenge: how to decide which potential foods to sample and which to pass by. Social species can address this challenge by distributing the risks of exploring the environment across multiple individuals. For instance, Norway rats share knowledge of food types and locations socially with other members of their communal burrow (Galef 1996). Rats will acquire a preference for a novel food smelt on the breath of a conspecific (Galef & Wigmore 1983), in spite of the rats' habitual neophobia. The key stimulus is the detection of carbon disulphide, a component of rat breath, in combination with a novel food smell (Galef *et al*. 1988). Clearly, this sort of social learning mechanism will be useful to opportunistic generalists like the Norway rat which exploit a wide range of food types and have frequent social encounters: far better to learn from others about what is safe to eat than to find out for oneself through costly trial and error.

However, paired with this ability to learn food preferences socially is a failure to learn aversions. Contrary to expectations, Galef *et al*. (1983) found that rats will acquire a preference for a novel food even if they smell it on the breath of a poisoned conspecific (see also Galef *et al*. 1990). If we frame the problem as learning what not to eat, then an animal might learn directly through observing the negative reactions of others or indirectly by copying the preferences of others and avoiding all else. Rats exhibit only this indirect social acquisition of aversions (Galef 1985).

*Author and address for correspondence: Informatics Research Institute, School of Computing, University of Leeds, Leeds LS2 9JT, UK (jasonn@comp.leeds.ac.uk).

†Present address: School of Cognitive and Computing Sciences, University of Sussex, Falmer, Brighton BN1 9QH, UK.

This failure to learn aversions is not due to an inability to perceive negative reactions to food in other rats. Experiments on the 'poisoned partner effect' (Lavin *et al*. 1980; Bond 1982) have shown that rats can in fact tell the difference between healthy and poisoned conspecifics. For example, Lavin *et al*. (1980) demonstrated taste aversion learning by rats for a liquid that they drank while caged with a poisoned partner. The presence of a sick rat is apparently a negative unconditioned stimulus.

Moreover, other opportunistic species are capable of direct social learning of aversions. Red-wing blackbirds (*Agelaius phoeniceus*) will learn to avoid an unpalatable food after a single observation of a conspecific's reaction to it (Mason 1988). Day-old chickens (*Gallus gallus*) have a similar ability (Johnston *et al*. 1998).

If rats can tell what others have been eating and can determine that another rat is sick, why are they unable to put these two pieces of information together to learn aversions socially, as blackbirds and chickens do? Galef (1985) suggested two related explanations. First, social learning is selected for given intermediate levels of temporal variation in the environment (Laland *et al*. 1996)—in highly homogeneous environments, genetically transmitted behavioural strategies will suffice and in highly varied environments, pure individual learning is favoured. Thus, it may be that the way in which novel toxins are distributed in the rat's environment happens not to favour the social learning of aversions. Second, it is possible that rats' individual adaptations for avoiding poisoning (e.g. neophobia and a dislike of bitter tastes) are so effective that there is minimal selection pressure for learning aversions socially.

We can assess the coherence of these explanations using an evolutionary simulation model in which individual rats are explicitly represented and their behavioural strategies permitted to evolve over time. This modelling

Received 7 June 2000 *Accepted* 14 September 2000

technique has been successfully employed in studies of the evolution of individual learning (e.g. Todd & Miller 1991; Belew & Mitchell 1996) and is increasingly popular in fields such as ecology (Judson 1994; Grimm 1999). We constructed a simulation in which a key feature of the environment, the lethality of toxins that might be ingested, was systematically varied and we looked at the effect of this variation on the evolution of discrimination between sick and healthy rats and its use in learning aversions socially. Our aim was to explore the match between particular ways that the environment can be structured and the social learning strategies best suited for those environments. In this way, we can in particular assess the possible explanations for the behaviour of Norway rats by identifying environmental conditions in which direct social learning of aversions is not selected for.

2. THE MODEL

In our model, a population of up to 500 simulated rats foraged from a cluster of 25 centrally located nests, exploiting 50 food sites. Each rat behaved in accordance with a genetically specified strategy allowing the potential for both individual and social learning. The cognitive and perceptual abilities of the rats had particular costs and error rates associated with them in order to match the uncertain world of real rats more closely. Rats that were successful in eating good food and avoiding toxins were more likely to reproduce. Simulation details and parameter values not described below are given in Appendix A.

(a) *The environment*

Parcels of food appeared at random sites in the environment at a constant rate. Most foods were (equally) nutritious, but 10% of food types were (equally) toxic. There was a certain probability that ingesting a toxic food type would result in a rat's death; we refer to this parameter as the lethality level. To ensure that the rats always had to deal with novelty, a list of 100 food types was used, but only a window of ten food types from the list could appear at any one time. Every 16 simulated days the window would advance one position up the list, so that one old food type stopped appearing and a novel one entered the scene. The life span of an individual rat (which was an emergent property of the simulation) never grew long enough for it to experience all 100 food types.

(b) *The rats*

Rats inherited four 'genes' that determined their learning strategy. Gene E determined how likely a rat was to eat novel food, gene S specified whether it would smell the breath of other rats in order to learn about what they had eaten, gene P gave the proportion of nest-mates that it would smell in this way and gene D specified whether it discriminated when smelling novel food on the breath of sick and healthy rats, allowing it to form aversions or preferences accordingly (without gene D, rats could only pick up preferences for novel foods smelt on the breath of others). Genes S and D were encoded as binary values (on or off), while genes E and P were represented as real numbers between 0.0 and 1.0. Gene E is an indicator of how much individual learning a rat is likely to perform (i.e. by trying new foods for itself), while genes S, P and D spell out a rat's social learning strategy.

In order to support the possibility of both preference and aversion learning, individual rats needed a memory for different food types. Any food a rat encountered was judged as either novel, preferred or aversive, depending on the rat's earlier experience of that food. For newborn rats, all foods were novel.

Each day, the rats foraged for food, interacted in the nest and possibly reproduced. When foraging, a rat decided whether or not to eat based on its memory for the food type it had found and, if that food was novel, on the rat's value for gene E (figure 1). There was a small probability of making an error in identifying foods before eating them: a preferred food might be confused with an aversive one or vice versa. However, novel foods were always recognized as such.

Eating good food provided energy and also meant that the rat would remember that food as preferred. Rats that had not been poisoned would return to the nest as soon as their stomachs were full, or, failing that, after five foraging periods had elapsed. A rat had to take in food in the short term in order to avoid starvation, as there was an energy cost for simply being alive. In the long term a rat needed to accumulate nutrients if it was to reproduce successfully.

Rats that ate toxins would experience either sickness or death; the probability of one or the other was determined by the lethality parameter associated with the environment. If a rat died it was immediately removed from the simulated world. Rats that became sick instead would develop an aversion for the food they had just eaten, return directly to the nest and show signs of sickness during one subsequent interaction period.

When interacting in the nest, rats could potentially smell the breath of other rats and learn either preferences alone or both preferences and aversions depending on their evolved strategies (figure 2). The smelling gene S specified whether a rat would smell the breath of its nest-mates and, thus, learn about what they had been eating. There was a small energy cost associated with expressing this gene, reflecting the costs of increased perceptual and behavioural sophistication.

The value of the P gene determined the proportion of nest-mates that a rat would attempt to smell if the S gene was active. Nest-mates were sampled randomly and with replacement so a rat with a high value of P would inevitably smell some nest-mates twice. We could have simply specified that a rat would smell a certain fraction of its possible partners, but given that so little is known about how real rats interact in their burrows, we chose to let this parameter evolve.

The discrimination gene D specified whether a rat would be sensitive to the state of health of a conspecific on whose breath it smelt a novel food, acquiring an aversion if the conspecific was sick and a preference otherwise. Thus, a rat with only the S gene would be capable solely of preference learning, whereas a rat with both S and D genes would also learn aversions. As with the smelling gene, there was a small energy cost for being able to

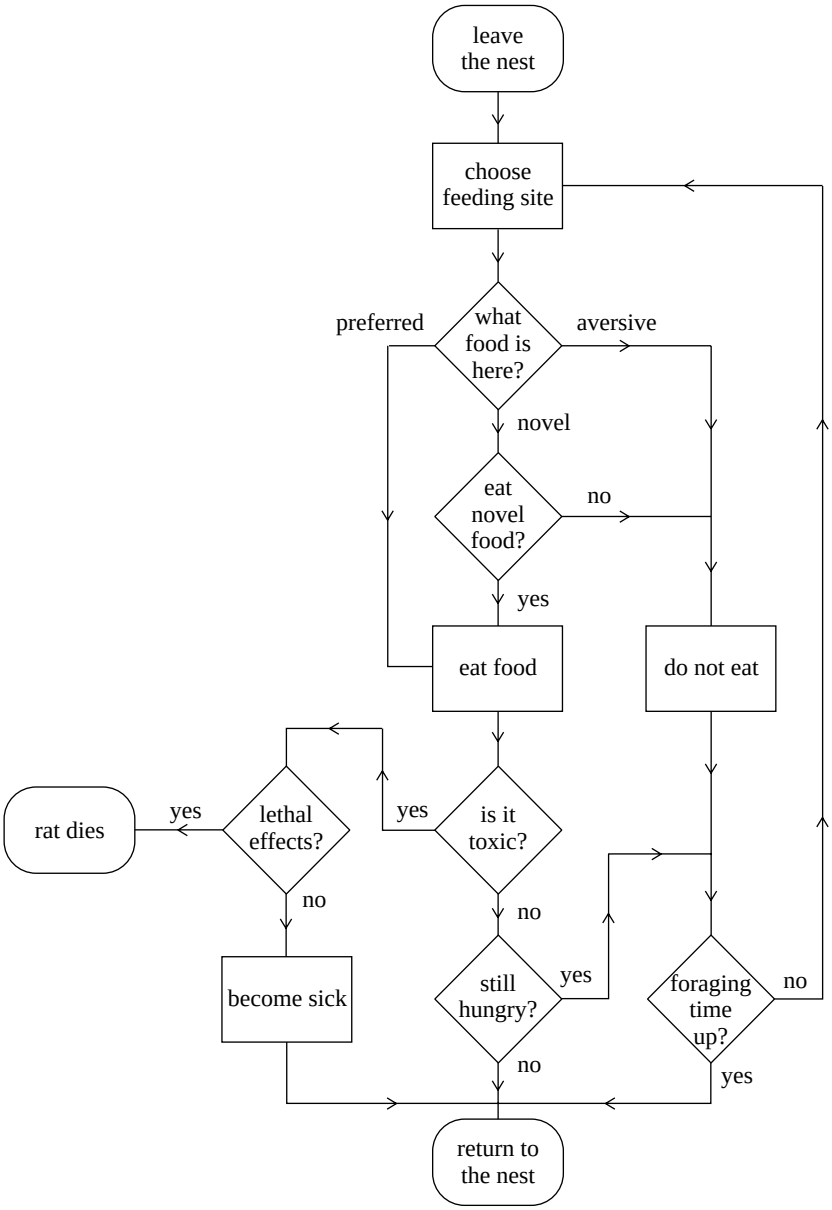

Figure 1. A flow chart showing the decision process for foraging rats. Each rat started by leaving the nest (top) and then chose a feeding site at random. If it found a preferred or aversive food, its choice was clear, whereas if it found a novel food the value for its E gene was used to determine randomly whether or not it would eat. If a food turned out to be toxic, a lethality parameter was used to determine randomly whether the rat died immediately or returned to the nest showing signs of sickness. Rats that did not eat or rats that ate a partial meal due to competition could select another foraging site, but after five such attempts their time was up.

discriminate. There was also a low probability that a rat would make a discrimination error in any one instance, mistaking a sick rat for a healthy one or vice versa.

A rat that accumulated a high level of food energy would reproduce asexually. A newborn rat inherited the four-gene behavioural strategy of its parent, with a small chance of mutation. The carrying capacity of the environment was fixed at 500. The population usually stayed below this level due to death by starvation and poisoning, but if necessary the current oldest rat would be killed to make room for a newborn. At birth, a reproductive cost was deducted from the parent rat's energy reserves and the newborn rat started with a store of energy somewhat lower than the cost of reproduction.

Initial populations in evolutionary simulations are often started with individuals having random genotypes. Pilot runs indicated that this would cause problems in our model, as random behaviour would often lead to population extinction in high-lethality environments. Our initial populations were therefore set up as conservative social learners: the S gene was present and the E gene was set to 0.005 in all rats, while the D and P genes were set randomly. Starting populations were also given accurate knowledge of the initial ten food types in the environment.

The choices of parameter values used in the model are not as constrained by empirical data as we would like. Although there is a great deal of information available on the behaviour of rats in the laboratory, data on the ecology of wild rats are not extensive (see e.g Lore & Flannelly 1997; Lore & Schultz 1989). We have therefore pitched the simulation at a relatively abstract level and,

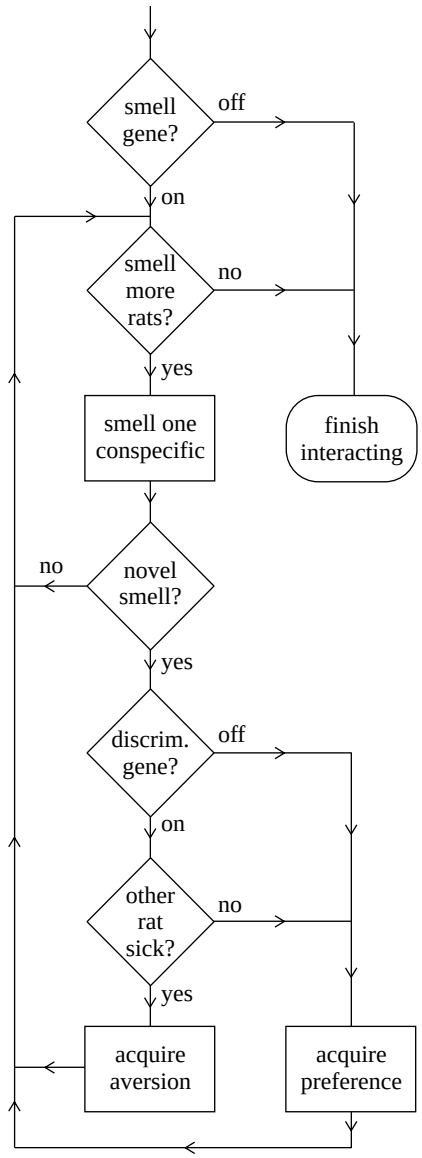

Figure 2. A flow chart showing the decision process for rats interacting in the nest. Rats without the S (smelling) gene would not interact at all. Rats with the S gene began by smelling the breath of a randomly chosen nest-mate. If they detected a novel food smell, their behaviour depended on the presence of the D (discrimination) gene. Rats with the D gene would examine the nest-mate's apparent health and learn either a preference or an aversion for the novel food smell accordingly. Rats without the D gene would acquire only a preference. The P gene specified how many times the process should be repeated with additional random nest-mates.

in so doing we hope to have captured some aspects of the selection pressures impinging on opportunistic social foragers in general.

3. SIMULATION RESULTS

The evolution of food aversion learning could be affected by several aspects of the simulated world we constructed: first, physiological features of the rats themselves, that is the noisiness of their perceptual and cognitive systems and the costs of running these, and second, the toxin-related features of the environment, specifically the frequency of toxin occurrence and the

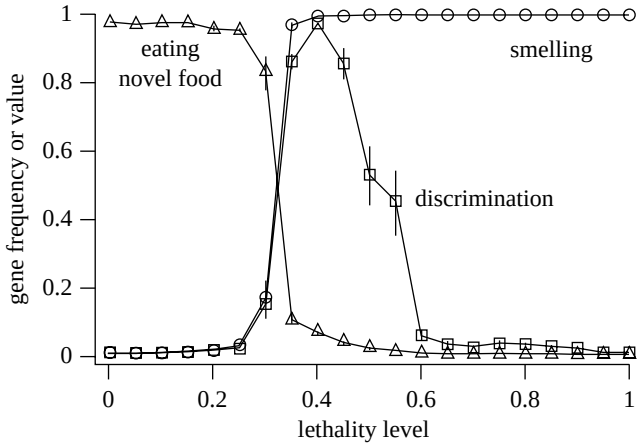

Figure 3. Evolved strategies by lethality level, the probability that eating toxic food would prove lethal. The lines for smelling and discrimination show the frequencies of the S and D genes, whereas the line for eating novel food shows the mean values of the E gene. When the lethality level was low, rats used individual learning (i.e. a high probability of eating novel foods). Faced with higher lethality levels, rats employed a social learning strategy (smelling each other's breath) but they only discriminated in a narrow middle range of lethality values. Thus, at high lethality levels, they matched the behaviour of real rats by socially learning preferences but not aversions. Each data point represents the mean state of the population after 25 000 days and error bars show the standard error across ten simulation runs.

chance of death associated with ingesting a toxin. Based on the results of earlier simulations, we fixed all but the last of these parameters at reasonable values. Because we expected toxin lethality to have the most impact on which learning strategies evolve, we varied just this important feature of the environment. We established that variation in costs and noise had no qualitative effect on the model's behaviour (see figure 7).

We carried out ten simulation runs for each of 21 levels of lethality in increments of 0.05 between 0.0 and 1.0 inclusive. Each run covered 25 000 simulated days, equivalent to over 80 generations of rats. The results we report show the mean state of the population at the end of the runs. Analyses of statistical significance have not been performed on the data, for two reasons. First, graphs of means and standard errors across runs clearly convey the outcomes of our model. Second, the reporting of statistical significance for simulation data is open to abuse, as arbitrary levels of significance can often be attained simply by running more simulations.

(a) *The basic model*

Figure 3 summarizes the evolved strategies in terms of the frequency of the E, S and D genes at different lethality levels. Three patterns are apparent. When lethality levels were low (i.e. in benign environments where eating something toxic was unlikely to kill the rat) individual learning was selected for. Rats almost always ate novel foods (high E gene value) and, thus, found out for themselves whether or not those foods were safe. The rats did not smell each other's breath (S gene usually off), indicating that social learning was not selected for under these conditions. (Discrimination is not relevant when

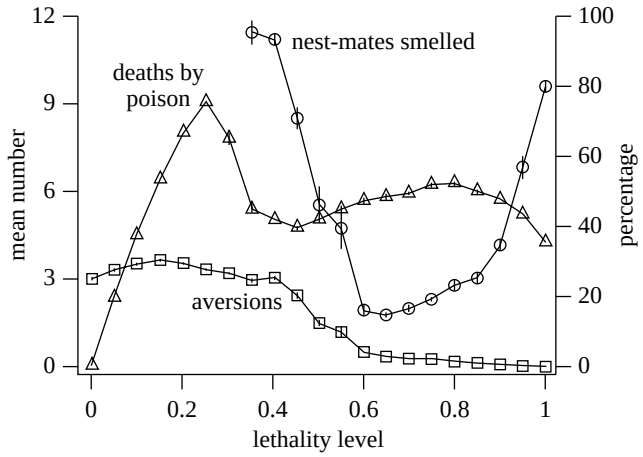

Figure 4. Population statistics by lethality level, the probability that eating toxic food would prove lethal. The number of nest-mates smelt (left axis) was calculated by multiplying the mean value of the P gene by the mean number of rats in a nest. This value is only plotted for lethality levels greater than 0.3 because at lower levels very few rats were smelling each other (figure 3) and the value of P was thus free to drift. The mean number of aversions held by each rat (left axis) falls off at higher lethality levels. The percentage of deaths caused by poisoning (right axis) rather than starvation or old age does not have a simple linear relationship with lethality, but peaks just before individual learning switches over to social learning. Each data point represents the mean state of the population after 25 000 days and error bars show the standard error across ten simulation runs.

the S gene is off, but the small discrimination cost kept the proportion of D genes low in this range.)

Between lethality levels of 0.3 and 0.35, there was an abrupt transition. For lethality values *ca.* 0.4, rats were much less likely to eat novel food. Individual learning declined in importance and social learning was selected for: both the smelling and discrimination genes approached fixation. The rats engaged in direct social learning of both preferences and aversions by interacting with others in the nest.

Faced with high lethality levels (>0.55) the rats were very unlikely to eat novel food. They continued to rely on social learning, but did not discriminate between sick and healthy conspecifics, acquiring a preference for any novel food smelt on the breath of another. This third pattern amounts to the indirect social learning of aversions; it matches the behaviour of real Norway rats and shows that discrimination between sick and healthy conspecifics may not always be selected for.

Figure 4 gives further details of the rats' behaviour and demography. When the rats were socially learning both aversions and preferences, they collected a lot of information, smelling about 11 of their nest-mates each day. As the lethality level increased, this figure dropped to as few as two, before increasing again in the most lethal environments. Clearly, a rat has to interact with at least one other rat if it is going to perform any social learning. However, whereas rats capable of direct social learning of aversions can safely interact with a large number of conspecifics, it seems that rats without this ability must limit their contacts in order to reduce the risk of acquiring a preference for a food that turns out to be toxic—at least until, at high lethality levels, the chance of encountering a poisoned rat has fallen sufficiently.

The mean number of aversions held by each rat declined as the lethality level increased (figure 4), falling off sharply with the transition from direct to indirect social aversion learning. It may seem counter-intuitive that aversions should become less frequent as toxins become more dangerous, but the result makes sense when we consider that a rat in a high-lethality environment almost never tries new food, is incapable of acquiring aversions socially and will only learn an aversion if it eats a toxic food and is lucky enough to survive. Figure 4 also shows that poisoning as a cause of death reached its peak value of *ca.* 80% just before the transition from individual to social learning. Social learning kicks in precisely when the risks of individual learning have become excessive.

(b) *Model variants*

We constructed variations on our basic simulation in order to explore possible explanations for these results. Our first concern was to understand why the discrimination gene, gene D, should be selected for at intermediate lethality levels but not at higher ones.

One possible answer had to do with the number of sick rats returning to the nest. In the basic model a rat that was fatally poisoned would die immediately. This meant that, as lethality levels increased, rats were less likely to carry information about toxic foods back to the nest. With lethality at 0.0, *ca.* 13 rats in every 1000 were showing signs of sickness on returning from foraging and, when lethality was equal to 0.4, and discrimination was at its height, the figure fell to 3.3 in every 1000. In the extreme case, with lethality equal to 1.0, no rats survived eating toxic food and there were never any sick rats to observe—under these circumstances, it is not surprising that there was no selection pressure for discrimination.

We therefore suspected that the absence of discrimination at high lethality levels could be due to the low incidence of sick rats returning to the nest. In other words, interacting with a sick rat, although potentially informative, becomes so rare that there is no selection pressure for distingishing between the sick and the healthy. In order to test this idea, we ran a simulation in which the frequency of sick rats returning to the nest was artificially maintained at a minimum of 3.3 in 1000. (Whenever the frequency dropped below this level, an extra rat was introduced into the simulation, made sick with a randomly chosen toxic food and placed into a random nest with one day to live.) Figure 5 shows that, in this case, discrimination was maintained even under the highest lethality levels. It follows that the rarity of interactions with sick conspecifics in the basic model is a good part of the causal story behind the fall off in discrimination we see in figure 3.

Another possible reason that the simulated rats did not discriminate at high lethality levels relates to the chance of making errors in discrimination. The function of discrimination must be to help the rat avoid developing a preference for a toxic food. However, given that discrimination is imperfect, this goal of correctly rejecting toxins must be balanced against the danger of falsely rejecting good foods. If an error-prone discrimination ability causes a rat to reject good food much more often than it

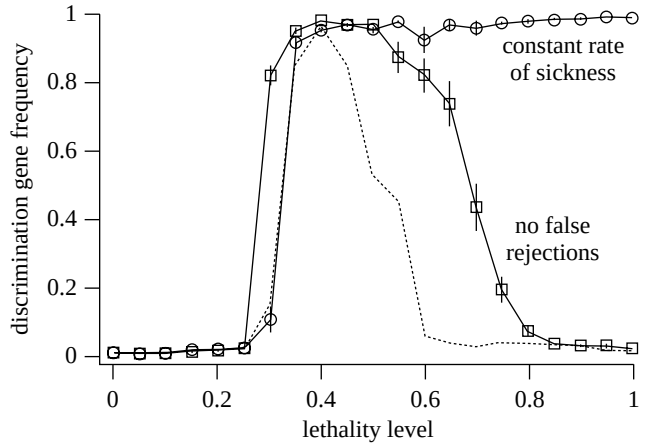

Figure 5. Frequency of the discrimination D gene by lethality level for two manipulations of the basic model—the dashed line shows the basic model's original frequency of discrimination (figure 3). When the number of sick rats returning to the nest was artificially maintained at a minimum of 3.3 in every 1000, discrimination was selected for right up to a lethality level of 1.0. When the rats were prevented from ever falsely rejecting a food type (i.e. forming an aversion to a good food) discrimination was selected for over a wider range of lethality levels than in the basic model. Both the rarity of sick rats and the costs of rejecting good foods thus contribute to the absence of discrimination at high lethality levels seen in the basic model. Each data point represents the mean state of the population after 25 000 days and error bars show the standard error across ten simulation runs.

saves the rat from poisoning, it may be better to stop trying to discriminate at all. In order to test this theory, we ran a second simulation variant in which discrimination errors were manipulated in such a way that the rats were prevented from ever falsely rejecting a good food. The logic was that, if rats discriminate more often in this variant, then the danger of false rejections must be partly behind the selection against discrimination at high lethality seen in the basic model. Figure 5 shows that selection for discrimination indeed occurred across a somewhat wider range of lethality values when errors were reduced. The need to avoid rejecting good food is thus one reason why the rats did not discriminate at high lethality levels, but it is a less important factor than the rarity of sick rats returning to the nest, as described above.

In the basic model, eating a toxic food resulted in immediate effects: either death or sickness. B. G. Galef (personal communication) noted that delays in the onset of sickness and prolonged periods of sickness before dying have been observed in wild rats. He suggested that these phenomena would probably have an effect on the rats' evolved learning strategies. In order to investigate this point we constructed several simulation variants in which delays were introduced (figure 6). First, we looked at what would happen if fatally poisoned rats were able to return to the nest and interact with others before dying. Figure 6 shows that under these circumstances, discrimination was selected for right up to a lethality level of 1.0. We also tried introducing random delays and durations for the effects of poisoning. In this variant the delay in days before becoming sick (previously zero), and the duration in days of sickness before the rat either died or recovered were randomly drawn from a normal distribution with a mean of 1.0 and a standard deviation of 1.0 (values were rounded to the nearest integer and negative values were set equal to zero). Figure 6 shows that, under these circumstances, there was some noise but discrimination was generally always selected for across high lethalities. Finally, we tried repeating the random delay variant, but specified that the duration of fatal sickness should be zero. In other words, there were delays before poisoning took effect, but only rats that were ultimately going to survive could experience an extended period of sickness. This means that any rat that had eaten a toxin and was now in the nest either appeared healthy (before the onset of toxic effects), in which case health-based discrimination would be erroneous or was sick but would recover, the same as in the basic model. In this case (figure 6), because discrimination could lead to more errors, it was even more strongly selected against than in the basic model. We conclude that a 'quick death' assumption is important in producing the fall off in discrimination shown in figure 3.

The basic version of our model incorporated energy costs for smelling and discrimination, as well as the possibility that rats would make errors in recalling food types and in discriminating between sick and healthy conspecifics. Empirical data on the magnitude of each of these factors are unavailable. We can safely assume that the real values are non-zero, but the values used in the simulation are somewhat arbitrary. In order to determine how important these costs and error rates were in producing the general pattern of our results, we re-ran the simulations with each factor set to zero in turn. The major aspects of the original model's behaviour (figure 3) were preserved: individual learning was always favoured at the lowest lethality levels, there was always a switch to social learning at some point as the lethality level increased, the smelling gene, gene S, would go to fixation for all lethality values above this point and selection for the discrimination gene, gene D, would increase and then fall away. However, figure 7 shows that manipulating the costs and error rates had a strong influence on the range of lethality values within which discrimination was favoured. The energy costs for smelling and discrimination clearly play a major role in producing selection for gene D over a narrow, intermediate range of lethality values. Memory and discrimination errors also have an effect on selection for gene D. For lethality values between 0.6 and 0.8, for example, both energy costs and both types of errors need to be in place for discrimination to be selected against.

4. DISCUSSION

The fact that Norway rats do not discriminate between sick and healthy conspecifics when learning socially about food may ultimately be explained by properties of the toxins in their environment. Our results show that this behaviour is surprising only when viewed from the perspective that gaining more information about the environment is always beneficial for an organism. However, here, as in many circumstances, the structure of the environment can simplify the individual's decision

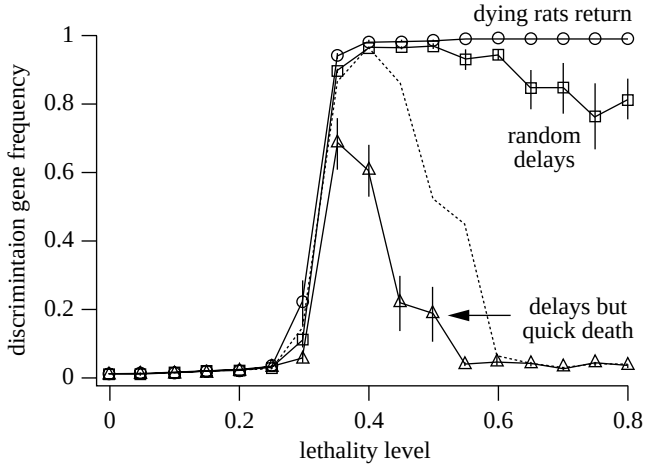

Figure 6. Frequency of the discrimination D gene by lethality level for variants of the basic model with time-delays introduced—the dashed line shows the basic model's original frequency of discrimination (figure 3). Data are presented only for lethality levels up to 0.8; at higher levels, time-delays made the environment extremely harsh and populations tended to become extinct. When fatally poisoned rats were able to return to the nest and exhibit sickness before dying rather than dying immediately, discrimination was selected for but not as strongly. With random delays before sickness, but a duration of zero for fatal sickness, discrimination was selected for even less often than in the basic model. Rapid death (without sickness) thus seems important in explaining the lack of direct aversion learning at high lethality levels. Each data point represents the mean state of the population after 25 000 days and error bars show the standard error across ten simulation runs.

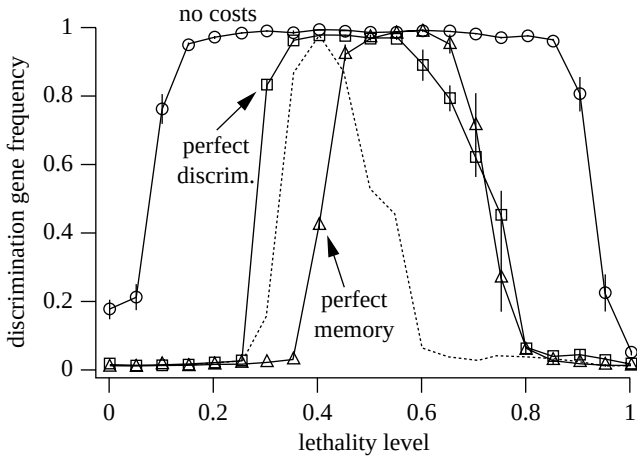

Figure 7. Frequency of the discrimination D gene by lethality level for three variants on the basic model—the dashed line shows the basic model's original frequency of discrimination (figure 3). With no energy costs for smelling or discrimination, the D gene was selected for over a much wider range of lethality values. With perfect discrimination (i.e. no chance of confusing sick with healthy conspecifics) the range of selection for the D gene also widened. When the rats' memories were perfect (i.e. no chance of confusing preferred with aversive food) the range in which the D gene was selected shifted to the right. Each data point represents the mean state of the population after 25 000 days and error bars show the standard error across ten simulation runs.

task, so that gathering further information is unnecessary (Gigerenzer *et al*. 1999). We hypothesize that rats do not learn aversions directly from each other because their social learning strategy evolved in an environment in which toxins were relatively lethal. This has led to a situation in which interacting with a sick rat—although it would provide useful information—does not happen often enough to make discrimination worthwhile. The advantages of direct social aversion learning may also be offset by the need to balance toxin avoidance with the costs of mistakenly rejecting good food.

Galef (1985) suggested that Norway rats fail to learn aversions socially because the distribution of toxins in their environment may happen not to favour this type of social learning and because their individual adaptations for avoiding toxins are so effective that there is no selection pressure for discrimination. Our results provide some support for both of these ideas: we have demonstrated that, in particular environments (i.e. high lethality ones) discrimination will not be selected for. Moreover, in these environments the combination of extreme neophobia and the social learning of preferences suffices for reproductive success and there would be no great advantage for a mutant that began discriminating. The model reminds us that 'the environment' includes not just external physical and chemical details such as toxin lethality, but also the social environment. Newborn rats in the simulation face an array of unfamiliar foods, some of which may kill them if ingested. However, they do not face this challenge alone—they are surrounded by more experienced conspecifics who are extremely conservative about trying new foods and whose preferences can therefore be safely copied.

Our work obviously has its limitations. The energy costs and error rates used in the simulation have an effect on the results, but we do not know how close these parameters are to their true values. Future modelling efforts might benefit from a more fine-grained treatment of time when looking at the issue of sickness delays and durations—in our simulation, time-periods shorter than one day could not be represented. Indeed, if we turn out to be seriously wrong at any of the points where we have had to make a reasonable guess in the absence of empirical data, then of course it is not clear that our results will hold. For example, if real rats have largely stable diets and only very occasionally encounter novel foods, then a different explanation for their failure to learn aversions might be required. However, our aim was not to construct a precise simulation that would result in quantitative predictions, but to explore a model with some heuristic value for the study of social learning.

A potential problem in our interpretation of the model is the finding that, when dying rats can return to the nest, discrimination is selected for right up to the highest lethality levels. If rats in the wild manage to return to their burrows after ingesting fatal amounts of a toxic substance and if they survive long enough to interact with other rats, then our model cannot be used to explain their failure to discriminate. (Note that this is the behaviour that rat poison manufacturers rely on—a slow death that allows other rats to pick up a preference for the poison—but whether or not this is how natural toxins usually work on rats is not known.) More data on how wild rats behave after ingesting a natural toxin would be useful.

We conclude with an empirical prediction arising from the model. The three basic patterns of evolved gene combinations observed across lethality levels in our results (see figure 3) suggest a prediction that foraging species will adopt one of three strategies depending on the characteristics of the toxins in their feeding environment. First, if a species enjoys very low levels of danger associated with eating toxic food, perhaps because of feeding specialization in a class of mostly safe foods or effective defences against poisons once ingested, these animals 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 their conspecifics. Blackbirds (Mason 1988) and chickens (Johnston *et al*. 1998) may fit this profile. Finally, in cases where the ingestion of a toxin is likely to prove fatal, 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 conspecifics. It may be that high levels of risk associated with toxin ingestion are characteristic of opportunistic generalist feeders such as the Norway rat. In order to test this prediction as it stands, we would need to measure the lethality of toxic foods in an animal's ancestral environment—this is a difficult problem. However, results from the simulation point to a cluster of other characteristics that can be expected to accompany each strategy. For example, the model predicts that blackbirds and chickens should maintain more aversions and interact with a greater fraction of their available conspecifics, than Norway rats. Empirical confirmation of these predictions would bolster our understanding of social information sharing and particularly why some species seem to care about the food-induced illness of their neighbours while others ignore it.

We would like to thank Seth Bullock, Martin Daly, Jennifer Davis, Bennett Galef, Ralph Hertwig, Andy Philippides, Peter Stuer, Bill Timberlake and two anonymous referees for their valuable comments and suggestions. E.T. was partly supported by studentship 2044801 from the Consiglio Nazionale delle Ricerche (Italy).

APPENDIX A: SIMULATION PARAMETER VALUES

Twenty-five parcels of food of randomly determined sizes ($\mu = 100$ units and $\sigma = 40$) appeared at random sites each day. The type of food was randomly selected from the ten available in the current food window and any food previously at that site was erased.

The probability of a rat making an error in identifying a food before eating it was normally 0.01.

The model incorporated two energy economies, implemented as a short-term and a long-term energy counter for each rat. A rat's short-term energy level had a starting value of 10 units and could never take on a higher value. There was an energy cost for simply being alive (1 unit day^{-1}), which meant that the maximum length of time a rat could survive without food was ten days. Eating good food provided 10 units of energy (for a full meal).

The cost of expressing the S and D genes was 0.2 energy units per gene per day. The probability that a rat would make a discrimination error, taking a sick rat to be a healthy one or vice versa, was 0.05. When interacting with another, a rat with the S gene would detect the food type previously eaten by its partner. If the partner had not eaten for more than three days, no food would be smelt.

A rat had to accumulate 1000 units of long-term energy in order to reproduce. At birth, a cost of 600 energy units was deducted from the parent rat's long-term energy counter, while the newborn started with 500 units of long-term energy.

The mutation rate was 0.01. The mutation operator for the S and D genes consisted of a simple bit-flip, whereas mutation for the E and P genes consisted of adding a random Gaussian term ($\mu = 0$ and $\sigma = 0.04$).

As a further defence against rapid extinctions, rats in the initial population were given a starting long-term energy value of 950 units, leaving them only 50 units to gain before reproducing.

REFERENCES

Belew, R. & Mitchell, M. (ed.) 1996 *Adaptive individuals in evolving populations: models and algorithms*. Reading, MA: Addison-Wesley.

Bond, N. W. 1982 Transferred odor aversions in adult rats. *Behav. Neural Biol.* **35**, 417–421.

Galef Jr, B. G. 1985 Direct and indirect behavioral pathways to the social transmission of food avoidance. *Annls NY Acad. Sci.* **443**, 203–215.

Galef Jr, B. G. 1996 Social enhancement of food preferences in Norway rats: a brief review. In *Social learning in animals: the roots of culture* (ed. C. M. Heyes & B. G. Galef Jr), pp. 49–64. San Diego, CA: Academic Press.

Galef Jr, B. G. & Wigmore, S. W. 1983 Transfer of information concerning distant foods: a laboratory investigation of the 'information centre' hypothesis. *Anim. Behav.* **31**, 748–758.

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*). *J. Comp. Psychol.* **97**, 358–363.

Galef Jr, B. G., Mason, J. R., Preti, G. & Bean, N. J. 1988 Carbon disulfide: a semiochemical mediating socially induced diet choice in rats. *Physiol. Behav.* **42**, 119–124.

Galef Jr, B. G., McQuoid, L. M. & Whiskin, E. E. 1990 Further evidence that Norway rats do not socially transmit learned aversions to toxic baits. *Anim. Learn. Behav.* **18**, 199–205.

Gigerenzer, G., Todd, P. M. & ABC Research Group 1999 *Simple heuristics that make us smart*. New York: Oxford University Press.

Grimm, V. 1999 Ten years of individual-based modelling in ecology: what have we learned and what could we learn in the future? *Ecol. Model.* **115**, 129–148.

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. *Anim. Behav.* **56**, 1347–1353.

Judson, O. P. 1994 The rise of the individual-based model in ecology. *Trends Ecol. Evol.* **9**, 9–14.

Laland, K. N., Richerson, P. J. & Boyd, R. 1996 Developing a theory of animal social learning. In *Social learning in animals: the roots of culture* (ed. C. M. Heyes & B. G. Galef Jr), pp. 129–154. San Diego, CA: Academic Press.

Lavin, M. J., Freise, B. & Coombes, S. 1980 Transferred flavor aversions in adult rats. *Behav. Neural Biol.* **28**, 25–33.

Lore, R. & Flannelly, K. 1977 Rat societies. *Sci. Am.* **236**, 106–116.

Lore, R. K. & Schultz, L. A. 1989 The ecology of wild rats: applications in the laboratory. In *Ethoexperimental approaches to the study of behavior*, vol. 48 (ed. R. J. Blanchard, P. F. Brain, D. C. Blanchard & S. Parmigiani), pp. 607–622. Boston, MD: Kluwer.

Mason, J. R. 1988 Direct and observational learning by redwinged blackbirds (*Agelaius phoeniceus*): the importance of complex visual stimuli. In *Social learning: psychological and biological perspectives* (ed. T. R. Zentall & B. G. Galef Jr), pp. 99–115. Hillsdale, NJ: Lawrence Erlbaum Associates.

Todd, P. M. & Miller, G. F. 1991 Exploring adaptive agency. II. Simulating the evolution of associative learning. In *From animals to animals: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, (ed. J.-A. Meyer & S. W. Wilson), pp. 306–315. Cambridge, MA: MIT Press/Bradford Books.

As this paper exceeds the maximum length normally permitted, the authors have agreed to contribute to production costs.