PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Learning and Memory in Mimicry. I. Simulations of Laboratory Experiments

John R. G. Turner and Michael P. Speed

Phil. Trans. R. Soc. Lond. B 1996 **351**, 1157-1170
doi: 10.1098/rstb.1996.0100

Learning and memory in mimicry. I. Simulations of laboratory experiments

JOHN R. G. TURNER[1]* AND MICHAEL P. SPEED[1,2]

[1] *Department of Genetics, University of Leeds, Leeds LS1 9JT, U.K.*
[2] *Department of Environmental and Biological Studies, Liverpool Hope University College, Stand Park Road, Childwall, Liverpool L16 9JD, U.K.*

CONTENTS

SUMMARY

An understanding of the dynamics of mimicry requires the modelling of the behaviour of predators in the wild. Current knowledge about the behavioural expression of learning and the dynamics of forgetting is insufficient for the construction of a definitive model of such behaviour. In particular, there is insufficient information on the response of a vertebrate when subjected to the presentation of a single identical conditioned stimulus paired with two (or more) unconditioned stimuli of different intensity or even of opposite effect, which in the present situation can be regarded as 'model' and 'mimic'.

A general algorithm of learning and forgetting, based on the behavioural model of Bush and Mosteller, is proposed; it is applied as a linear operator in Monte Carlo simulations of the behaviour of a predator confronting a mixture of models and mimics. The algorithm is varied in detail: learning may be cumulative or instantaneous, constant or variable according to the strength of stimulus, and towards a continuously distributed or two-state (0 and 1) asymptote; forgetting may be cumulative or instantaneous, constant or variable according to the strength of stimulus, and dependent on time or on the occurrence of external events. Thirty different behaviour systems arise from rational combinations of the various learning and forgetting rules, including as special cases those behavioural models already proposed in the literature on mimicry.

A standard experimental technique is the presentation to predators of a constant number of prey with varying proportions of models and mimics: we term this a reciprocal frequency experiment, and simulate its outcome for all thirty rational models. The results of such experiments, when appropriately transformed, will yield straight lines or curves according to the behavioural model employed. Models with all-or-none features (instantaneous learning or forgetting) tend to yield straight lines: curves tend to appear when the model assumes gradual or cumulative learning and forgetting. The result is dominated by the mode of learning; forgetting plays a secondary part. The method will therefore discriminate well between 'switched' and 'gradual' modes of learning and forgetting, but only if the experiments are carried out over a wide range of palatabilities for both the model and the mimic, and with other adequate design features. It is also necessary to design the experiment to distinguish between the dynamics of short- and long-term learning, which may well be different.

---

* Present address: Department of Biology, University of Leeds, Leeds LS1 9JT, U.K.

Experiments to date permit no definitive or even tentative choice between the behavioural models, and do not necessarily, contrary to the literature, unambiguously support the models of discontinuous learning and forgetting proposed by J. E. Huheey. Experiments by M. L. Avery do, however, suggest tentatively that the dynamics of the short-term assessment of strategy by birds confronted with an intimate mixture of models and mimics may be adequately modelled as a process of rapid all-or-none switching between high and low attack probabilities, at least if the model and mimic are strongly aversive and strongly appetitive. It appears that there is a major gap in the understanding not only of mimicry systems but of the dynamics of learning and forgetting as motivational drivers in vertebrates.

1. INTRODUCTION

Mimicry is the result of coevolution between two species, termed the model and the mimic, usually driven by the behaviour of a third species, the 'signal receiver' or 'dupe' (Wickler 1968; Vane-Wright 1976; Pasteur 1982); in the most-studied forms of mimicry, this third species is a predator, which drives the coevolution between the colour patterns of two or more species of potential prey (Gilbert 1983; Turner 1984*a*, *b*, 1995). The evolutionary and ecological dynamics of this kind of mimicry is thus controlled by the psychology of the predator.

It may therefore come as something of a surprise, after mimicry has been intensely studied for over a century (Bates 1862; Kimler 1983*a*, *b*) and conditioning for over half a century (Pavlov 1927), to find that we do not yet know enough about predator psychology to describe with any certainty the dynamics of mimicry systems in the wild. As behavioural ecologists we are interested in precise descriptions of behavioural change and therefore define learning in predators as changes in the probability of attack. The interest of contemporary animal psychology, with its emphasis on mechanisms of knowledge acquisition (see, for example, Mackintosh 1983) rather than motivation and behavioural production has given surprisingly little data on which to model predation in the real world. Specifically, there is a shortage of psychological research into the nature of learning when the same conditioned stimulus may ambiguously represent diverse forms of the unconditioned stimulus, for instance either a very pleasant or a very unpleasant experience. This ambiguity is the essence of mimicry, and the paucity of research into it exposes not only our ignorance of mimicry but a serious gap in our understanding of the dynamics of learning: despite considerable advances in research into the acquisition of information, there are no clear, simple or obvious rules for determining the relation between sensory input and muscular output. In addition, although acquisition or 'learning' has been much researched, psychologists seem not to clearly understand the characteristics of subsequent forgetting or mind-changing in natural or quasi-natural settings.

Thus convention assumes that Batesian mimicry, i.e. mimicry of an unpalatable species by a palatable prey, benefits the mimic but harms the model. If the mimic population expands, then the advantage enjoyed by the mimic is reduced and the harm to the model increases (see, for example, Ford 1975; Sheppard 1975). If on the other hand the mimic is unpalatable then both model and mimic should gain protection, particularly as the mimic population grows. This type of mimicry is termed Muellerian. The evolutionary dynamics of these two types of mimicry are thought to differ considerably (Turner 1977, 1987). This view of mimicry has been challenged, especially by Huheey (1976, 1980*a*, 1988), Sbordoni *et al.* (1979), Rothschild (1980, 1981), Owen & Owen (1984) and Speed (1993). Rationally constructed but *a priori* models of predator behaviour suggested that the division of mimicry into Batesian and Muellerian types was simplistic, and either that some of the properties of these two types of mimicry were not as convention supposed, or that the categories were inappropriate, or indeed that it might be better to recognize quite new categories. Perhaps all mimicry could be seen as 'essentially' Batesian, or 'essentially' Muellerian. On the other hand, other rational arguments based on different *a priori* assumptions about predator behaviour (Benson 1977; Sheppard and Turner 1977; Turner *et al.* 1984; Turner 1984*a*, *b*) supported the conventional view of mimicry.

The sometimes heated debate has so far suffered from an inadequate exploration of the models of predator behaviour on which the various arguments were based: all participants have tended to make unanalysed assumptions about behaviour. The argument has further ignored models of learning developed by psychologists (Bush & Mosteller 1955; Rescorla & Wagner 1972; Mackintosh 1975; Pearce & Hall 1980; see Speed 1993).

In this pair of papers we therefore: (1) propose a generalized, simple behavioural model of learning behaviour, which includes all the models so far proposed as special cases, and further extend it to embrace a variety of other models not so far explored; (2) examine the results expected from each of these models in the types of learning experiment (reciprocal frequency experiments (see, for example, Huheey 1988)) which have been used to investigate the behaviour of predators confronted with mixtures of models and mimics; and (3) (in our second paper) examine the dynamics to which a range of these behavioural models would lead in the wild.

We conclude that the present experimental evidence is inconclusive, and that therefore the dynamics of mimicry in the wild remain unknown.

2. BEHAVIOUR SYSTEMS

(*a*) *Predator psychology*

We use a simple Monte Carlo simulation of the behaviour of a vertebrate predator (Turner *et al.* 1984). The controlling variable that governs its response to prey is a fluctuating probability of attacking any particular category of prey. This probability is decreased if the predator eats an unpalatable prey, and likewise increases after a palatable prey is eaten; indeed, these are clearly operational definitions of palatable and unpalatable prey. Thus the consumption of an unpalatable prey will result in a period of avoidance, during which the same prey species may be observed but tends not to be attacked. At the end of such a period of avoidance the probability of attack reverts to some higher value, and therefore the predator reverts to attacking the prey again, having either forgotten its unpalatability or made a decision to probe the environment. In the absence of further reinforcing experiences, the probability will, in this way, sooner or later reach an asymptotic value at which the predator could be described as being in a completely naive state. Similarly, periods of repeated attacks on palatable prey may be terminated by the probability declining towards a lower level. The predator's hunger level and other motivational states remain constant throughout, reflecting the usual experimental situation in which the predator is kept adequately fed with an alternative food supply, or a situation in the wild where alternative food is available and the predator is able to attack the model-mimic system opportunistically. All simulations are commenced with the predator in the naive state with regard to all prey.

Such a model of behaviour can be varied in (1) the asymptotic value of this attack probability after an infinite number of encounters, (2) the rate of approach to the asymptote, (3) the nature of the forgetting process, and (4) the naive attack probability (although in practice we have held this constant). These processes of learning and forgetting are described separately below in the form of rules. Different combinations of rules constitute what we will term behaviour systems (or behavioural models: the term 'model' is not used further, for obvious reasons).

(*b*) *Ecology and predator behaviour*

The predator is imagined as encountering prey of varying kinds during a series of time intervals; it is useful to think of it as a sit-and-wait predator like a jacamar (see, for example, Chai 1986) watching butterflies fly past, or as a foraging predator which meets a sequence of prey during its search, or as a caged predator presented with prey by the experimenter in a regular timed sequence. In any one time interval either one or no prey individuals are encountered.

The predator encounters a sequence of models and mimics at random in all time intervals, in simulations of laboratory experiments (this paper); or four species of prey in a random sequence and at frequencies proportional to their imagined densities, in simulations of mimicry systems in the wild (Speed & Turner 1996; and as in Turner *et al.* 1984). *Model* is warningly coloured and uniformly noxious to the predator. *Mimic* is identical to Model in its external appearance but can vary in flavour right through the spectrum of palatabilities; that is, the predator always reacts in the same way to Model and Mimic on sight alone, but the consequences of sampling by the predator are different for the two types of prey. The two further species will be used only in our second paper.

The general characteristics of simulated behaviour are readily described: in any one run of the experiment, the predator starts with its probabilities of attack in the naive state, and consequently sooner or later attacks any prey presented. It rapidly learns to avoid an unpalatable prey, subjecting it to only intermittent attacks, rapidly develops a liking for palatable prey, and exhibits a more complicated sequence of avoidance and attack on a Model–Mimic pair (sample sequence in Turner *et al.* 1984). The results were recorded as percentages of the prey attacked, over a sequence of time intervals (normally 8000, after probabilities of attack are allowed to equilibrate during a sequence of 1000 time intervals that are not monitored). Presentation of the prey, and whether or not the predator attacks, are determined by the standard Monte Carlo method of comparing the prey density or attack probability with a pseudorandom number. Simulations were carried out on a Dell 433 s/L PC using Turbo Pascal.

(*c*) *Learning rules*

We simulate learning as a post-attack modification of attack probability. We use five learning models, all of which can be described as cases of the general algorithm of Bush & Mosteller (1955) (see also Wagner & Rescorla 1972):

$$P_2 = P_1 + \alpha(\lambda - P_1), \qquad (1)$$

where P_1 is the probability of attack at the start of an encounter and P_2 is the attack probability after the attack. The variable α is the learning-rate variable, which can take any value within the limits 1 (single-trial learning) and 0 (no learning); λ is the learning asymptote for attack probabilities, and its limits are 0 and 1. In all models, λ takes some value in the range $0 \leqslant \lambda < 0.5$ for unpalatable prey and $0.5 < \lambda \leqslant 1$ for palatable prey; at the point of neutral palatability λ is 0.5. The naive attack probability P_0 is set at 0.5 for all predators, even where published models have assumed otherwise (in their original forms the rules of Huheey (1964) and Owen & Owen (1984) set it to 1).

This equation is a simple linear operator which, upon iteration, produces a decelerating change in attack probabilities; it is applied in any time interval in which the predator attacks a prey individual. Clearly in a time interval when there is no attack, $P_2 = P_1$. It can be seen that each prey encounter decreases or increases the attack probability by a fraction α of its current distance from the asymptote λ. With palatable prey, the attack probability generally increases; with unpalatable prey it generally decreases. It may help to note that Bush & Mosteller (1955) and Turner *et al.* (1984) used the equation in the form

$$P_2 = \lambda + \beta(P_1 - \lambda) \text{ or } P_2 = \beta P_1 + (1-\beta)\lambda, \qquad (2)$$

where $\beta = 1 - \alpha$.

Table 1. *General properties of the five learning rules*

rule	mode	α (rate variable)	asymptote distribution
GR	gradual	variable	continuous
BM	gradual	constant	continuous
ST	gradual	variable	discontinuous 0,1
OO	instantaneous[a]	[constant]	continuous
HH	instantaneous[a]	[constant]	discontinuous 0,0.5,1

[a] or 'switched' hence α constant.

The learning rules differ in the procedures by which the values of α and λ are determined, or in general terms in the mode of learning, the constancy or variability of the rate variable, and the distribution of the asymptotes (table 1).

(i) *Rule 1: Generalized learning rule (GR)*

This expands the Bush and Mosteller learning rule to its most general form. The variable λ takes any value, according to the degree of palatability, in the continuous range from and including 0 to 1. For neutrally palatable prey, λ is 0.5; its value is therefore $0.5 < \lambda \leqslant 1$ for palatable prey (with 1 the most palatable), and $0 \leqslant \lambda < 0.5$ for unpalatable prey (again with 0 being extremely unpalatable). The value of α varies with intensity of prey palatability from 1 to 0.5 (most to least intense), according to the formula

$$\alpha = 0.5 + |\lambda - 0.5|.$$

Therefore in this rule both the rate of learning and the asymptote of learning are dependent on prey palatability. Learning is instantaneous with highly palatable or thoroughly unpalatable prey and more gradual or cumulative with milder prey.

(ii) *Rule 2: Bush and Mosteller learning rule (BM)*

In this application of Bush and Mosteller's model (Bush & Mosteller 1955) the rate of learning is constant but its asymptote is variable: that is, control of the value of λ is the same as in the generalized learning rule (GR) but α is set at 0.5 for all prey. Therefore, although the asymptote of the attack probability varies, the rate of approach to this asymptote resulting from attacks is constant; learning is always a gradual or cumulative process.

(iii) *Rule 3: Sheppard and Turner learning rule (ST)*

This rule was devised by J. R. G. Turner in 1962 (unpublished; effectively an independent reinvention of the Bush–Mosteller model), developed later by P. M. Sheppard (unpublished) and finally described by Turner *et al.* (1984); formally, it is a variant of the generalized model (GR) in which the asymptote of learning is constant and the rate of learning variable (the inverse of the Bush and Mosteller rule). Here λ has only two possible values: 1 for all palatable prey, and 0 for all unpalatable prey. The value of α depends upon the degree of prey palatability or unpalatability: for the very highly palatable or unpalatable its value is 1, for the neutrally palatable it is 0, and for those with palatabilities of intermediate intensity α is between these values. Learning is therefore complete in one trial with extremely palatable or unpalatable prey, non-existent with neutrally palatable prey and of intermediate rate with more mildly flavoured prey. With prey of different degrees of palatability or unpalatability, given a sufficient number and rate of encounters, the attack probability always reaches 1 or 0, respectively.

(iv) *Rule 4: Owen and Owen learning rule (OO)*

This can be classed as a variant of the Bush and Mosteller rule (BM) with instantaneous learning. The value of λ is determined in the same way as in the GR and BM learning rules: the asymptote of learning can take any value from 0 to 1 according to palatability. Here α is constant with a value of 1; learning is thus always completed in a single trial (Owen & Owen 1984).

(v) *Rule 5: Huheey learning rule (HH)*

This model (Huheey 1964, 1976) generates immediate switching of the predator's behaviour after sampling a prey item, and can be classed as a discontinuous version of the OO rule. As in the ST learning rule λ is 1 for all palatable prey and 0 for all unpalatable prey. For simulation purposes, we have introduced a further value, not used by Huheey, of 0.5 for prey of exactly neutral palatability. In all cases α is 1, as in the OO rule: learning is instantaneous, and attack always leaves the attack probability at either 1, 0 or 0.5 depending on the class of prey palatability. This rule and the OO rule can both be solved explicitly when combined with appropriate forgetting rules (Huheey 1964, 1976; Owen & Owen 1984), but for the sake of uniformity we have tested them by simulation.

(*d*) *Forgetting rules*

Forgetting is simulated as a process which counteracts the effects of learning by returning disturbed attack probabilities to the naive value. We use a simple algorithm of the Bush and Mosteller type

$$P_3 = P_2 + \phi(P_0 - P_2), \tag{3}$$

where P_2 and P_3 are the probabilities of attack before and after forgetting, and ϕ is a variable that controls the rate of forgetting. P_0 is the naive attack probability and identically the asymptote of the forgetting process; its value is 0.5 in all our simulations, even in those behaviour systems whose original authors set it to 1 (Huheey 1964, 1976; Owen & Owen 1984). In the absence of learning, equation (3) returns the attack probability asymptotically to its naive value.

In all behaviour systems this operator is used on all attack probabilities at least after each time interval, regardless of whether learning has occurred within that interval. The rules that define the values of ϕ vary considerably in the mode of forgetting, the variability of the rate variable and whether forgetting is triggered by time and/or the passage of specific events (table 2).

Table 2. *General properties of the seven forgetting rules*

rule	mode	ϕ (rate variable)	trigger
ST	gradual	constant	time
OO	gradual	variable	time
HH	instantaneous[a]	[constant]	events
HT	instantaneous[a]	[constant]	time
HS	gradual	constant	events
CO	gradual	constant	time and events
OP	gradual	constant	time negated by events

[a] or 'switched' hence ϕ constant when activated.

(i) *Rule 1: Sheppard and Turner forgetting rule (ST)*

In this rule ϕ is always 0.02, causing a two percent reversion of attack probabilities towards 0.5 at the end of each time interval (Turner *et al.* 1984).

(ii) *Rule 2: Owen and Owen forgetting rule (OO)*

Here ϕ varies between 0 and 0.05 according to the intensity of prey palatability (most to least intense: from no forgetting to rather rapid forgetting). This is an adaptation of Owen and Owen's assumption that the intensity of prey palatability experienced at an encounter determines the rate of forgetting about that encounter (Owen & Owen 1984).

(iii) *Rule 3: Huheey counting and forgetting rule (HH)*

In this rule ϕ is fixed at 0 after an attack (i.e. the predator does not forget the experience) until a certain number of apparently identical prey (that is, the same species, or its mimic, or its model) have been observed and – by definition – avoided. The value of ϕ then becomes 1; that is, the predator immediately forgets and the attack probability returns to the naive state (Huheey 1964, 1976). The number n of such avoided prey is a measure of the unpalatability of the victim that was sampled; for neutrally palatable prey the number is 0, whereas for unpalatable prey its value is large. Because this predator never subsequently avoids sampled palatable prey, it never has the opportunity to forget experiences with them; thus in this case ϕ always remains at 0.

Our introduction of a value of 0.5 for the naive attack probability adds a stochastic element to this rule, absent from Huheey's original formulation. Huheey's predator avoids n prey and reverts to attack; our predator avoids n prey and then avoids each subsequent prey with a probability of 0.5 until an attack occurs. Our avoidance number therefore has a minimum value n plus a short stochastic 'tail'.

(iv) *Rule 4: Huheey time-based rule (HT)*

This is our own variation of Huheey's (1964) rule (HH: rule 3), which substitutes the passage of time for the number of prey observed (in some of Huheey's original applications these two aspects of the process were not operationally distinguished, as it was a consequence of the particular conditions that the number of prey observed in a standard time interval was constant; see also Luedeman *et al.* (1981)). Again ϕ takes values of 0 (after an attack) and 1 (when conditions for forgetting are fulfilled) but ϕ is 0 until a specified length of time has elapsed. The duration of this period of remembering, which is unchanged by events of seeing and avoiding prey, reflects the intensity of prey palatability; it is short for mildly unpalatable prey and long for those that are highly unpalatable. Because this mechanism is not dependent on avoiding prey, it operates for palatable as well as unpalatable prey, and as with the HH rule (above) has a stochastic element. Luedeman *et al.* (1981) give some explicit solutions for a rule similar to this one, in which the passage of time is measured by the regular appearance of prey of a number of different species.

(v) *Rule 5: Forgetting through avoiding: a gradual version of Huheey's rule (HS)*

This is our own gradual (and therefore fully stochastic) version of Huheey's rule (HH: rule 3). In this case ϕ is 0 for all attack probabilities unless an identical prey is observed and avoided when it is 0.02: forgetting is thus event-based, as in Huheey's rule, but occurs cumulatively rather than as a form of switching. This is based on an implication of Huheey's equations (Huheey 1964, 1976) that seeing and avoiding a prey is the only source of forgetting.

(vi) *Rule 6: Forgetting depends on time and avoiding: a combined rule (CO)*

This rule is a combination of the ST forgetting rule (rule 2) and the HS rule (rule 5); ϕ is 0.02 throughout. The forgetting routine is activated for all attack probabilities at the end of each time interval and *additionally* for attack probabilities on the same species or its model or its mimic when identical individuals are seen but not attacked.

(vii) *Rule 7: Observing-prevents-forgetting rule (OP)*

This rule is a modification of the ST forgetting rule (rule 1): ϕ is 0.02 for every time interval unless an identical prey is observed, when ϕ for the appropriate attack probability is 0. Observing a prey therefore temporarily halts forgetting ('reminds' the predator) and this rule holds whether a prey is attacked or not.

(*e*) **Describing behaviour systems**

Permutations of learning and forgetting rules can be assembled into thirty-five behaviour systems, thirty of which we have used. Each behaviour system is identified according to its learning and forgetting rules. For instance, the behaviour system proposed by Sheppard and Turner (Turner *et al.* 1984) will be referred to as ST-ST, because it uses those learning and forgetting rules respectively. Huheey's behaviour system (Huheey 1964, 1976) was HH-HH, and Owen & Owen's (1984) is OO-OO. A behaviour system that combined Sheppard and Turner learning with Huheey forgetting would be ST-HH, and so on. The 'cognitive' behaviour system explored by Speed (1993) is, in these terms, GR-ST.

There are a number of other descriptions of learning (see, for example, Rescorla & Wagner 1972; Mackintosh 1975; Pearce & Hall 1980) and suggestions for modelling forgetting (see, for example, MacNamara & Houston 1987), which we do not consider here but which may give the same patterns as some of our behaviour systems. As a test of behaviour systems our method is therefore indicative rather than decisive; simple behavioural models of the Bush and Mosteller type have the advantage that they can readily be explicitly simulated.

3. EXPERIMENTAL TESTS OF BEHAVIOUR SYSTEMS BY HUHEEY'S RECIPROCAL FREQUENCY METHOD

(*a*) *Background*

In laboratory investigations of mimicry it is customary to present a constant number of prey to a predator and to vary the proportion of models and their Batesian mimics (Brower 1960; Huheey 1980*b*; Avery 1985). Thus one might present 50 prey in a fixed number of days with some predators receiving 10 models and 40 mimics, and others receiving 30 models and 20 mimics. We can categorize this as a *reciprocal frequency experiment*. This type of experiment has been used to validate an important behaviour system. Huheey (1964) showed that under these conditions his behaviour system (HH-HH) predicts a linear relation of the form

$$(1/P)-p = nq, \tag{4}$$

where P is the proportion of presented models and mimics (combined) that the predator attacks, and p is the percentage frequency of mimics presented; n is the number of prey avoided after an attack, and $q(=1-p)$ is the frequency of models. A plot of $1/P-p$ against q therefore produces a straight line with gradient n. Experiments on amphibian (Huheey 1980*b*) and avian predators (Brower 1960; Avery 1985) have all produced such straight lines, and therefore seem to verify Huheey's models of learning and forgetting (Huheey 1988).

To find out whether this type of experiment will in fact serve as a critical test for Huheey's behaviour system, we have carried out simulations of reciprocal frequency experiments based on thirty rational behaviour systems with various combinations of learning and forgetting rules (see table 5).

(*b*) *Simulation procedures*

For these experiments prey were presented in all time intervals, as a random sequence of Model and Mimic. Thirty behaviour systems were tested. The data from each simulation have been plotted (figure 1), following Huheey (1964, 1980*b*, 1988), in the form $1/P-p$ against $p(=1-q)$; we have noted whether this produces a straight line or a curve. Each graph point is the result of 8 replications of runs of 8000 cycles after an equilibration run of 1000 cycles (with the exception of simulations with ST-OP, to avoid an infinite value of $1/P$). Error bars are two standard errors on either side of the mean. In view of this large number of trials, the straightness of a line was judged by eye rather than statistically (if the edge of a ruler would lie within all error bars of the plot then the line was taken to be straight). Each behaviour system was used in four palatability combinations, as follows.

Table 3. *Parameter settings for all learning rules in the simulations of reciprocal frequency experiments*

(For α and λ see equation (1). In the HH learning rule, models of different unpalatability are distinguished by n, the number of prey avoided during forgetting (table 4); mimics of different palatability are not distinguished.)

		value			
		mimic palatability		model unpalatability	
learning rule(s)	defining parameter	mild	high	mild	high
GR, BM, OO	λ	0.51	0.99	0.49	0.01
ST, HH	λ	1	1	0	0
GR	α	0.51	0.99	0.51	0.99
BM	α	0.5	0.5	0.5	0.5
ST	α	0.01	0.99	0.01	0.99
OO, HH	α	1	1	1	1

1. Combination 1: very unpalatable Model, very palatable Mimic.
2. Combination 2: mildly unpalatable Model, very palatable Mimic.
3. Combination 3: mildly unpalatable Model, mildly palatable Mimic.
4. Combination 4: very unpalatable Model, mildly palatable Mimic.

Table 3 shows the parameter settings for these combinations, for all learning rules.

Some of the forgetting rules determine their own parameter settings, but to combine certain learning and forgetting rules we have used a set of algorithms. When the HH or HT forgetting rule is combined with the GR, BM, or OO learning rule, the avoid number (that is, the number seen and avoided or the number of elapsed time intervals before a switch in forgetting occurs) for each prey is given by

$$n = 1+\mathrm{INT}\{20(|\lambda-0.5|)\}, \tag{5}$$

where INT indicates that only the integral part is retained. Thus if the prey were such that $\lambda = 0.01$, then the number avoided would be 10. For ST-HH and ST-HT the avoid number is given by

$$n = 1+\mathrm{INT}\{10\alpha\}.$$

When, for example, Model is extremely unpalatable ($\alpha = 0.99$), avoid number = 10 prey (or avoid time = 10 intervals). For systems using the OO forgetting rule with GR, BM, ST or OO learning, the rate of forgetting is given by

$$\phi = x/y,$$

Table 4. *Parameter settings for all forgetting rules in the simulations of reciprocal frequency experiments*

(For n (the number avoided during forgetting) see equation (5); for ϕ see equation (3).)

		value			
	defining	mimic palatability		model unpalatability	
forgetting rule(s)	parameter	mild	high	mild	high
ST, CO	ϕ	0.02	0.02	0.02	0.02
OO[a]	ϕ	0.049	0.001	0.049	0.001
OO[b]	ϕ	0.0495	0.0005	0.0495	0.0005
HS, OP	ϕ[c]	0 or 0.02		0 or 0.02	
HH, HT[d]	n	–	–	1	10
HH, HT[e]	n	1	10	1	10

[a] with GR, BM and OO learning rules; [b] with ST learning rule; [c] see rules for application in text; [d] with HH learning rule (mimics of different palatabilities are not distinguished); [e] with other learning rules.

where $x = 0.5 - |0.5 - \lambda|$, $y = 10$ for the GR, BM and OO learning rules, and $x = 1 - \alpha$, $y = 20$ for the ST learning rule. Table 4 shows the resulting settings for the forgetting parameters.

(*c*) *Results*

We have run many, but not all possible, behaviour systems. Several are not reasonable combinations (for example, a predator described by HH-ST would make no operational distinction at all between prey of differing unpalatability). Figure 1 shows examples of the graphs produced by these simulations; table 5 summarizes the full results, classified according to whether they produce a detectable curve (and if so, whether it is concave up or convex up), a straight line (hereafter line for short), or a result of the type shown in figure 1(*d*), with a line with very high standard errors (for a full set of figures see Speed (1990), or copies filed with the Royal Society editorial office). The table further marks those results that are shown to be curved in our 8000 replications, but which are so close to being linear that they would be completely indistinguishable from straight lines even in quite extensive experiments. It can be seen that the result is far from straightforward, but the following generalizations can be made.

Almost without exception, combination 3 produces a line (the exceptions are ST-HS and ST-OP): as Model and Mimic are of almost equal palatability in this case (just to either side of neutrality) it is unlikely that anyone would knowingly perform this experiment, which is of theoretical interest only. Therefore, omitting this combination from further consideration (that is noting behaviour in combinations 1, 2 and 4 only) we find that certain learning and forgetting rules are largely consistent in their performance.

Learning rules. Two rules, OO and HH, are consistent in producing lines (with very high errors for OO-OP in combinations 1 and 4). BM is consistently curved (except for combination 2 with HH and HT), and these are all concave upwards. GR produces a majority of curves (fourteen against seven lines) and these are mostly concave; however, three of the curves are effectively indistinguishable from lines. ST gives a majority of curves (twelve: one being very like a straight line), seven lines, and two 'complicated' results, the curves all being concave in combinations 1 and 2, and convex in combination 4. Overall, GR and ST are rather similar. *No rule consistently gives a curve*, although BM comes close: one is likely to get a straight line *under some experimental conditions* with any of the rules. The rules that consistently give lines (OO and HH) have the common property of instantaneous learning. The GR and ST rules (except when combined with OO or OP forgetting) also give straight lines, or curves that are not distinguishable from lines, when both the model and the mimic are powerful stimuli (very unpleasant and very pleasant): in these circumstances these two rules also give instantaneous learning.

Forgetting rules. Under the conditions of the simulation, in which a prey is presented in each time interval, HH (which counts the prey) and HT (which counts the time intervals) are formally identical, and give the same results: ignoring combinations with the OO and HH learning rules, which always give lines, these two forgetting rules give nearly equal numbers of curves (always concave) and lines (4:5). All other forgetting rules predominantly produce curves, unless they are overridden by the overwhelming straight-line effect of the OO learning rule. Ignoring combinations with this learning rule (and with HH), HS divides three lines to six curves (five of which are concave); OO has a single exceptional straight line, and one curve that is very close; CO has two lines; OP is either curved or 'complicated'; and ST is always curved, but has three curves that are virtually indistinguishable from lines. Among the curves, the proportions convex:concave are OO 0:8, ST 2:7, HS 1:5, CO 2:5; thus all rules are predominantly concave. None of the forgetting rules gives a result as consistent as the OO and HH learning rules, but again, even discounting the lines produced by combination with the OO and HH learning rules, all forgetting rules give a line or something very close with *some* combination of learning rule and palatability. The OP rule is the exception,

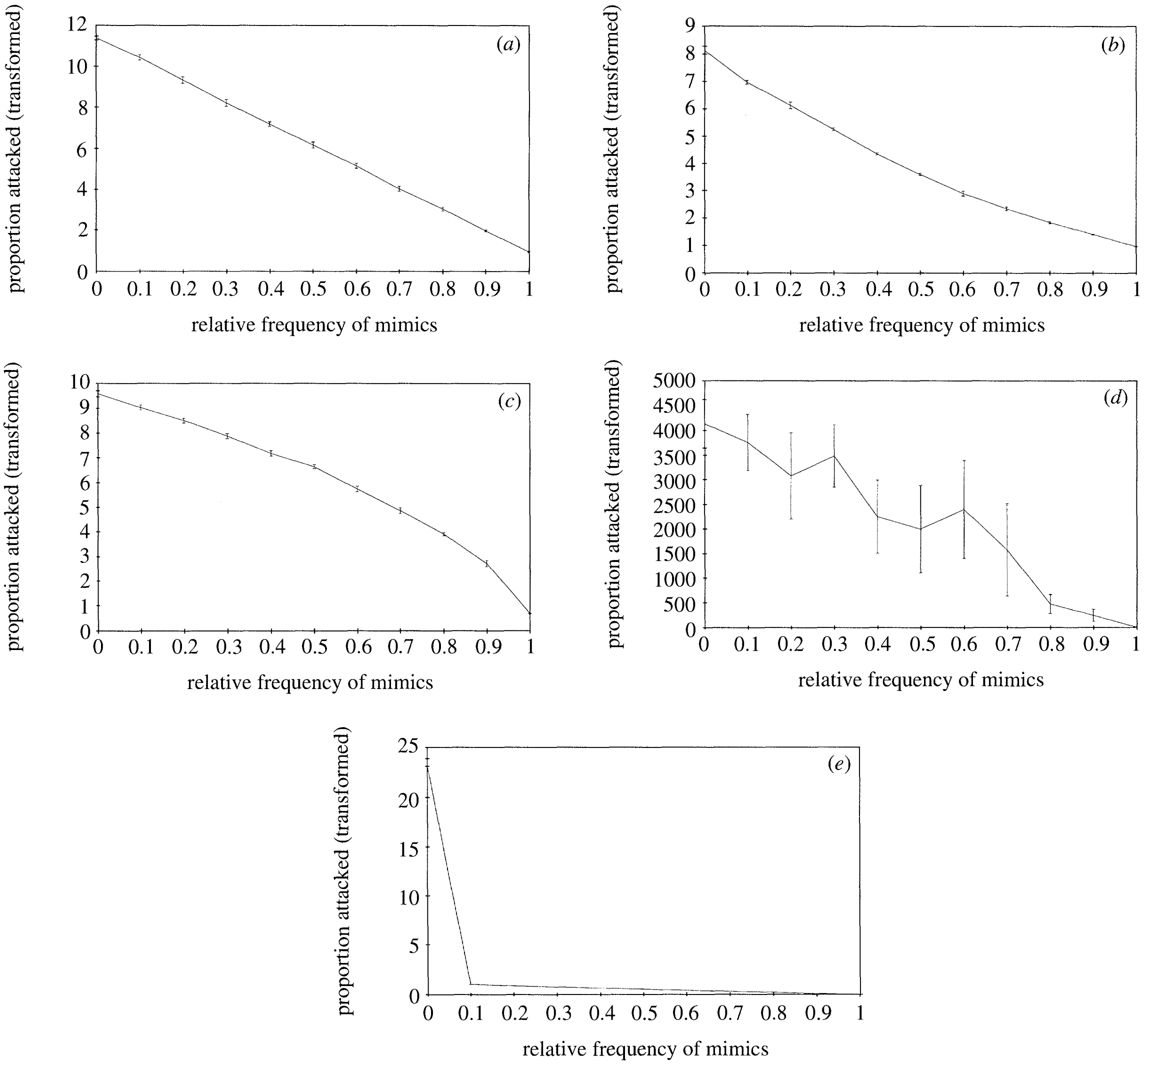

Figure 1. Examples of plots of simulations of reciprocal frequency experiments, with the proportion attacked transformed according to equation (4), plotted against the relative frequency of mimics in the mixture. All plots show data for combination 4, except (*e*), which is for combination 2. (*a*) A straight line (OO-CO); (*b*) concave upward (BM-ST); (*c*) convex upward (ST-CO), (*d*) a 'complicated' result with high values of $1/P-p$ and large errors (ST-OP), (*e*) a curve that is sharply elbowed (ST-OP). Error bars are two standard errors on either side of the mean.

Table 5. *Outcomes of reciprocal frequency simulations*

Symbols: –, straight line; ∪, curve, concave upward; ∩, curve, convex upward; !, irregular line with very high errors and extremely high values of $1/P-p$; •, not investigated.

Palatability combinations (see text) are tabulated in the order 1 2
4 3.

Rules that involve all-or-none switching are printed in **bold**; GR and ST learning switch when the Model and Mimic are strongly aversive and appetitive (combination 1).

learning rule	Forgetting rule													
	ST		OO		**HH**		**HT**		HS		CO		OP	
GR	$\cup^{a}$	∪	∪	∪	–	∪	–	∪	–	∪	–	∪	∪	∪
	$\cap^{a}$	–	$\cup^{a}$	–	–	–	–	–	–	–	∩	–	∪	–
BM	∪	∪	∪	∪	∪	–	∪	–	∪	∪	∪	∪	∪	∪
	∪	–	∪	–	∪	–	∪	–	∪	–	∪	–	∪	–
ST	$\cup^{a}$	$\cup^{c}$	∪	$\cup^{c}$	–	$\cup^{c}$	–	$\cup^{c}$	–	$\cup^{c}$	–	$\cup^{c}$	!	$\cup^{c}$
	∩	–	–	–	–	–	–	–	∩	$\cap^{c}$	∩	–	!	∪
OO	–	–	–	–	–	–	–	–	–	–	–	–	$\_^{b}$	–
	–	–	–	–	–	–	–	–	–	–	–	–	$\_^{b}$	–
HH	•	•	•	•	–	–	–	–	•	•	•	•	•	•
	•	•	•	•	•	•	–	–	•	•	•	•	•	•

[a] Virtually indistinguishable from a straight line; [b] very high standard errors; [c] 'elbowed' (figure 1*e*).

producing a simple line only with the OO learning rule, and otherwise producing curves or 'complicated' lines with extremely high values of $1/P-p$ (figure 1*d*). These results with ST-OP are generated without the 1000 equilibration runs, because by the end of such a period the attack probability has become permanently zero. This occurs because in this context with total prey density at one hundred percent, the OP rule prevents any forgetting if there is a highly unpalatable Model in the system (combinations 1 and 4). Some of the irregularity of these results therefore results from non-equilibrium behaviour. The rules producing the largest proportion of lines are HH and HT, which involve instantaneous switching of the motivational state: all other rules involve more or less gradual loss of information.

Behaviour systems. Again ignoring combination 3 (near-neutrality for both prey types), some complete systems (combinations of learning and forgetting) are consistently linear or curved. Systems that can be expected to produce lines under all circumstances are all systems with OO learning and all the tested systems with HH learning. Systems that are always curved are all systems with BM learning rules (except when combined with HH or HT forgetting), GR-OP and (curved but not necessarily effectively distinguishable from lines in real experiments) GR-ST, GR-OO and ST-ST. These systems should therefore give a consistent result (either linear or curved) provided at least one of the prey is strongly flavoured (either a highly palatable mimic or a highly unpalatable model), as must be the case for most experiments. The ST and GR learning rules when combined with HH or HT forgetting produce straight lines if the Model is highly unpalatable. The remaining systems give a mixture of curves or lines according to the palatability combination; these are GR-HS, GR-CO, ST-OO, ST-HS and ST-CO. These systems are clearly compatible with any experimental result, unless the two palatabilities are carefully monitored.

There are two further problems in recognizing lines and curves. First, there are lines that are sharply 'elbowed', being effectively indistinguishable from a straight line across much of the range of frequencies, and turning sharply at frequencies of zero or 1 (figure 1*e*), or becoming steeply curved over about half the frequency range: with our parameter settings these are most notably combination 2 with ST learning and any forgetting model, together with an exceptionally steep curve with ST-HS in combination 3. If only part of the frequency range were investigated, these could be mistaken for lines. Some further cases (notably ST-ST-4, ST-OP-3, BM-HH-1, BM-HT-1: the number is the combination) that are steep exponential curves could also be confused with lines over the flatter part of the range. Second, it is likely that when, for instance, as in the ST-ST system, combination 2 produces a concave curve and combination 4 a convex one, then (as Model and Mimic are both at their extremes – mild or strong – in the two combinations), experiments with moderately palatable mimics and moderately unpalatable models (as well as the very mildly flavoured ones in combination 3) might produce straight lines. We consider that it would constitute theoretical overkill to simulate this situation, which inspection reveals to appear four (or five) times in the table in transitions from combination 2 to combination 4; thus table 5 is if anything conservative in depicting the potential for producing straight lines in reciprocal frequency experiments.

General pattern. The overall pattern is that straight lines are consistently produced if the system contains the OO learning rule or the HH learning rule. Straight lines are also strongly produced with the HH and HT forgetting rules. Curves are usually produced otherwise, overwhelmingly by the BM learning rule (with one exception), and predominantly by the GR and ST learning rules. The OP forgetting rule generates exceptionally complicated results, with very high values of $1/P-p$ with the ST and OO learning rules and very unpalatable models. The mixture of concave and convex curves does not have a strong overall pattern, but convexity is characteristic of the ST learning rule (and to a lesser extent of GR) in combination 4; the BM learning rule is always concave. Over all systems convexity is much in the minority.

Therefore systems in which behaviour 'switches' between fixed motivational states, to attack or to avoid, produce straight lines in reciprocal frequency experiments. This is true even if the forgetting process involves a gradual loss of information or the ability to act on it, provided that learning is 'switched' (i.e. instantaneous in one trial); the effect persists, although less strongly, when learning is gradual and cumulative and it is forgetting that is 'switched'. Experiments giving lines over a range of palatabilities of model and mimic therefore indicate switching behaviour: this is shown by the HH and OO learning systems, which are consistently switched, and by GR and ST when the prey are so strongly appetitive and aversive as to cause these learning rules to switch. Curves, on the other hand, indicate that learning and forgetting are gradual (discounting cases where the curves are very close to straight, or where a concave and convex curve are likely to give a line at intermediate palatability for model, or mimic, or both); the BM learning system, which is consistently gradual, always gives a curve unless forgetting is switched (that is except BM-HH and BM-HT); GR-OP is also consistently curved. These curves should always be concave. All other systems are likely to give mixtures of lines and curves as the palatabilities are varied. (All the above of course excludes the lines given when both the model and mimic are near to neutrality.) All learning rules and all forgetting rules can give lines under some circumstances; only experiments conducted with a wide range of palatabilities will give clear results.

4. DISCUSSION

Simulations based on each palatability combination separate the behaviour systems into three classes when the results are plotted according to equation (4) (ignoring experiments with the nearly-neutral prey in combination 3): those that produce a straight line with all three combinations of palatability, those that produce a curve with all combinations, and those that produce a mixture of lines and curves. Reciprocal frequency experiments, if performed with a wide enough range of palatability of both the model and the mimic, will therefore distinguish quite well between behaviour systems, particularly in telling us, from the predominance of straight lines or curves, whether learning and forgetting are switched or cumulative. The type of result given by these experiments is dominated by the mode of learning (whether it is instantaneous or gradual) and there is a lesser influence of the forgetting mode (similarly): rapid switching, particularly instantaneous, single-trial learning, tends to produce a straight line. A single result, however, conducted with only one palatability of the model and the mimic, will usually be equivocal, as most systems will produce a straight line under some circumstances. A definite curve, however, would eliminate any kind of 'switched' or instantaneous learning; the system most nearly eliminated by a single straight-line result is the BM learning rule combined with gradual (i.e. not HH or HT) forgetting.

Huheey (1988, pp. 28–29) has suggested that reciprocal frequency experiments conducted with live predators have provided support for his model of predator learning and forgetting (HH-HH in this paper), by generating straight lines in the plot of $1/P-p$ against q. Our simulations show that, contrary to this expectation, the experiment does not provide any such critical confirmation unless it is carried out over a wide range of palatabilities. Although the HH-HH rule does indeed generate a straight line, in a limited experiment it is indistinguishable in this respect from many other behaviour systems. Straight lines are indeed produced by switched behaviour – instantaneous learning and instantaneous forgetting (whether time- or event-based) – all of which are indeed combined in the HH-HH system. However, all other systems that contain these features tend to produce straight lines as well. Conversely although gradual learning and forgetting tend to produce curves, they unfortunately do not invariably do so; a great many combinations of learning and forgetting produce straight lines with at least some combinations of palatabilities.

Empirically, therefore, the existing experiments (Brower 1960; Huheey 1980*b*; Avery 1985) do not allow us to draw any firm conclusions about the linearity of the relationship. An adequate experiment would be replicated so as to permit a statistical test of departure from linearity; it would explore the full range of model:mimic ratios, from $q = 0$ to $q = 1$, because some of the 'elbowed' and exponential curves will not be readily distinguishable from straight lines if only part of this range is investigated. The experiment should further be repeated with different strengths of palatability for the mimic and different degrees of unpalatability for the model. It needs also to be designed in such a way that short- and long-term learning effects are distinguished, as these may use very different classes of neural representation (see, for example, Wagner 1981) and must operate in different ecological contexts. In natural situations long-term memory will be more usually relevant to questions of mimicry; short-term memory will be in use only in the rather unusual case that the predator encounters a swarm of models with the mimics intimately mixed into it. This situation can occur with some rain-forest butterflies (Bates 1862; Brown & Benson 1974).

On these criteria all the existing experiments are flawed in one way or another. None has explored different palatabilities. The experiments of Avery (1985) were conducted in such a way as to involve a very large element of short-term memory mixed with longer-term learning: the birds were confronted for one hour at a time, at intervals of twenty-four hours, by a bowl of seeds containing an intimate mixture of treated (models) and untreated (mimics) seeds. Brower's experiments (Brower 1960) similarly involved the presentation of usually 10 replications (one experimental mealworm along with a control) in rapid succession, separated by 24 h between trials, and must also involve a large element of short-term memory. Only Huheey's (1980*b*) experiments are designed to test longer-term memory alone, with at least three hours between each presentation. Unfortunately although they were replicated the full data are not reported; hence these experiments cannot at present be further analysed, and we must rely on a subjective impression that the points, distributed inadequately (only in the range of q from 0.5 to 1) lie in a line.

It is likely, from the description of the aversive stimulus (quinine, methiocarb or a sting: table 2 of Huheey (1988)) that all the experiments used a strongly aversive model, and therefore should yield results in our combinations 1 or 4. As the mimic was 'normal' food, it is reasonable to suppose that it was highly palatable, and that the experiments therefore mostly involve our combination 1 (highly unpalatable model, highly palatable mimic).

Table 6 analyses Avery's experiments and, with a somewhat creative fusion of her two experiments, those of Brower for deviations from linearity. Neither of Avery's experiments shows a significant deviation of the regression of $1/P-p$ on q from a straight line: removal of an anomalous sextet of birds used to test the highest frequency of models in the second experiment produces an even closer fit to a straight line (experiment II: truncated). It appears that Avery's experiments do confirm Huheey's prediction that the relation between the two variables is linear. Brower's experiment does not depart significantly from a straight line, but is inconclusive as the non-significant variance between groups accompanied by the highly significant linear regression indicates bad behaviour of the error variance. Curing this problem, by removal of the points where there is only one replicate, results in all the sources of variation becoming non-significant (not tabulated).

Avery's two experiments therefore fail to show a departure from a straight line and hence present a weak *prima facie* case for supposing that learning which uses short-term memory or immediate strategy assessment can legitimately be modelled by a system involving rapid all-or-none switching between high and low attack probabilities, depending on the nature

Table 6. *Analysis of regression on the memory experiments of Avery (1985) and Brower (1960)*

(The analysis tests the regression of $1/P-p$ against q (proportion of models) as the independent variable. See text for symbols. **Avery experiments:** P is the percentage attacked out of total prey presented, but as the experiment is not designed in such a way that this last figure can be determined (the birds controlled their own approach to the feeding bowls), P is entered as the fresh mass (grams) of seed consumed in one hour, divided by the mean mass over the whole experiment of control ('baseline') seed consumed in one hour; respectively 1.44, 0.54, 0.54 g. This untreated control seed, presented simultaneously with the experimental mixtures of models and mimics, has been used also to represent consumption with $q = 0$. The other values of q are 0.25, 0.5, 0.75, and 1.0. In neither experiment is there a significant overall deviation from a straight line; in experiment II, truncation by the removal of the $q = 1$ data, which relate to a different group of birds, makes the results even more homogeneously linear. Deviation from linearity is invariant with the scaling of P, and F values for other sources of variation are comparatively robust to reasonable changes in the value of the divisor. The scaling does permit an estimate of n, the mean scaled mass of seeds avoided after an encounter with a model, directly from the regression coefficients: respectively for experiments I, II and II truncated, 15.3, 9.1 and 11.3 g. **Brower experiments:** P is the proportion of mimics eaten or attacked, as attacks on models are not tabulated. The values of q (each with one recorded value of P) in the two experiments are respectively 0, 0.1, 0.7 and 0.9 and 0.1, 0.4, 0.7 and 0.9. The experiments have been treated as replicates even though they differed slightly in the experimental conditions; they are analysed as the number of prey actually eaten or as the number attacked in any way. The values of n are respectively 13.1 and 14.9.)

source	sum of squares	d.f.	mean square	F	*Prob.*
Avery, experiment I					
between levels	916.36	4	229.09	12.7	< 0.001
linear regression	833.00	1	833.00	30.0	< 0.025
deviations from linear	**83.36**	**3**	**27.79**	**1.6**	**> 0.05**
within levels	413.63	23	17.98		
total	1329.99	27			
Avery, experiment II: in full					
between levels	337.68	4	84.42	28.42	< 0.001
linear regression	313.95	1	313.95	39.70	< 0.01
deviations from linear	**23.73**	**3**	**7.91**	**2.7**	**> 0.05**
within levels	74.25	25	2.97		
total	411.93	29			
Avery, experiment II: truncated					
between levels	277.47	3	92.49	29.7	< 0.001
linear regression	277.34	1	277.34	4217.1	< 0.001
deviations from linear	**0.13**	**2**	**0.07**	**0.02**	**no test**
within levels	71.72	23	3.12		
total	349.19	26			
Brower: number eaten					
between levels	171.80	4	42.95	2.9	> 0.1
linear regression	166.37	1	166.37	91.9	< 0.005
deviations from linear	**0.13**	**3**	**0.07**	**0.02**	**no test**
within levels	45.01	3	15.00		
total	216.81	7			
Brower: number attacked					
between levels	225.55	4	56.39	0.9	no test
linear regression	217.44	1	217.44	80.4	< 0.005
deviations from linear	**8.11**	**3**	**2.70**	**0.04**	**no test**
within levels	188.06	3	62.69		
total	413.61	7			

of the most recently sampled prey, with in this case an encounter with a model inducing a period of avoidance in which an average of around 12 g (scaled) of seed or 14 mealworms (avoid number n: see table 6 headnote) is eaten from the control offering before the bird returns to the model–mimic mixture. Thus if this is considered to be a convincing straight line produced in combination 1, it eliminates all systems with thoroughly gradual learning – that is all those using the BM rule – as well as a miscellaneous set of systems, namely GR-OO, GR-OP, ST-OO and ST-OP. However, it allows not only those systems with the OO or HH fully 'switched' learning rules to stand, but also the systems with the GR and ST learning rules, which are in general 'gradual', but which 'switch' when presented with prey in combination 1. Unfortunately, therefore, although the result demonstrates *prima facie* that the experimental birds were indulging in rapidly 'switched' changes in behaviour, they tell us nothing as to whether this switching is a general feature of learning and forgetting (or strategy assessment) within this time scale, or whether it is simply the result of the model and mimic being strongly aversive and appetitive. The result is compatible with systems with diametrically opposite assumptions: those that always switch and those that are in general gradual and cumulative. If on the other hand the palatabilities cannot be reliably supposed to be in combination 1, the production of a straight line does not then eliminate any of the learning rules or any of the forgetting rules.

Clearly the experiments fail to give even this limited certainty for the dynamics of learning involving long-term memory. All the experiments require to be repeated in such a way as to separate clearly the effects of long- and short-term learning, with full replication under a range of palatabilities, and over the full range of frequencies.

Thus the experiments taken at face value not only fail to eliminate any learning or any forgetting rule, but among specific behaviour systems fail to eliminate not only the original HH-HH system that they have appeared to support, but also a wide variety of others, including some like GR-ST whose assumptions differ radically from those of HH-HH. Thus HH-HH assumes instantaneous learning, extreme asymptotes of learning and switched, event-based forgetting, where GR-ST assumes cumulative learning, variable asymptotes of learning, and forgetting that is cumulative with the passage of time; only systems like BM learning, which can never be provoked into switching, are apparently eliminated.

One further system probably can be eliminated: under the conditions of this investigation, those behaviour systems that incorporate the observing-prevents-forgetting rule (i.e. all learning rules combined with the OP rule) never forget about a Model or a Mimic because, with the total Model–Mimic population at saturation throughout, Model or Mimic is seen in every time interval. With the ST-OP behaviour system the extremely high and unconvincing values of $1/P-p$ (maximum around 4000) when Model is very unpleasant suggest that this behaviour system can be quickly rejected; Brower's (1960) experiments used a very unpleasant prey and yielded values of the order of tens.

5. CONCLUSIONS

In behaviourist terms, the things which we need to know about learning and forgetting in order to produce models that will give satisfactory descriptions of natural mimicry systems are as follows (tables 1 and 2):

1. Asymptote of learning: extreme values (0,1) or a continuous variable depending on intensity of stimulus?
2. Learning: instantaneous or cumulative?
3. Learning rate variable: constant or a continuous variable dependent on intensity of stimulus?
4. Forgetting: time-dependent or event-dependent?
5. Forgetting: instantaneous or cumulative?
6. Forgetting rate: constant or variable according to intensity of stimulus?

It would be rather easy to investigate some of these points: for instance, if forgetting is dependent on time in the sense that it is produced by thermodynamic molecular changes in the brain, or is triggered by other mental activities, then it should be temperature-dependent, a point readily investigated by using amphibians, reptiles or other ectotherms (see Gleitman (1971) for goldfish).

It is beyond our scope at present to suggest extensive experimental designs: what is clear is that the evidence we have at present arises from reciprocal frequency experiments. Although these are by no means perfect discriminators between various behavioural systems (table 5), the consistent production of straight lines would indicate that learning and probably forgetting were single-trial, 'switched', instantaneous effects; the general production of curves would indicate that learning and forgetting were gradual and cumulative. It has been suggested (Huheey 1988) that the three sets of experiments to date (Brower 1960; Huheey 1980*b*; Avery 1985) support the former suggestion. We believe that such a conclusion would cause some scepticism among experimental psychologists but unfortunately (or fortunately, according to one's prejudices) these experiments have been performed with barely adequate replication, with an insufficient range of palatabilities of model and mimic, and in one case with an insufficient range of frequencies of model and mimic; thus they fail to make any clear discrimination between the behaviour systems on offer. Two of the experiments tested an inseparable mixture of long- and short-term learning (or strategy formation); the best of them at the most indicates that short-term strategy in the face of strongly aversive and appetitive prey (but not necessarily in general) involves rapid switching. Far from giving adequate support to the behavioural models of Huheey (1964, 1976, 1980*a*, 1988), current reciprocal frequency experiments leave the field wide open for the acceptance of almost any behavioural system.

Mimicry is so widespread (see, for example, Wickler 1968) that encounters with models and mimics must be an everyday occurrence for many vertebrates. Yet the emphasis in modern animal psychology on the mechanisms of knowledge acquisition (see, for example, Rescorla & Wagner 1972; Pearce & Hall 1980; Mackintosh 1983) at the expense of its long-term retention (but see Bouton 1994) and of the integration of memory with motivation, makes modelling difficult. This is true even in the apparently simple case of Muellerian or Batesian mimicry, whose the essence is the presentation of an almost identical conditioned stimulus with a conditioning stimulus that varies in intensity between encounters, or may even completely reverse its meaning from appetitive to aversive. Mimicry aficionados for their part have largely ignored the literature on the psychology of learning (few, for instance, have cited the work of Bush and Mosteller). Yet all predators at least must adapt to this situation (there is indeed a small literature on what an 'optimal' or 'well-designed' predator 'ought' to do: Estabrook & Jespersen (1974), Bobisud & Potratz (1976), Luedeman *et al.* (1981)); our knowledge of learning is severely incomplete until their response to it is understood. Likewise the dynamics of mimicry will not be comprehended.

In general, the fact that we cannot declare any learning–forgetting model canonical, and then apply it to the dynamics of mimicry, shows that there is a serious gap in the understanding of the psychology of learning: specifically of the rules which relate the acquisition of knowledge and its retention to motivation and to overt behaviour.

We are grateful to J. Wright of the Psychology Department in the University of Leeds for much general advice on the relevant literature in the psychology of learning, and for reading the draft of this paper; to A. Goodwin for IT support; and to Dr. Andrew T. D. Bennett of the Animal Behaviour Unit Department, University of Bristol, and to two anonymous referees for some very constructive suggestions. M. P. Speed has been supported by the SERC and the Research Committee of Liverpool Hope University College, and J. R. G. Turner by the Academic Development Fund of the University of Leeds and by an equipment grant from the Royal Society.

REFERENCES

Avery, M. L. 1985 Applications of mimicry theory to bird damage control. *J. Wildl. Mgmt* **49**, 116–1121.

Bates, H. W. 1862 Contributions to an insect fauna of the Amazon Valley. Lepidoptera: Heliconidae. *Trans. Linn. Soc.Lond.* **23**, 495–566.

Benson, W. W. 1977 On the supposed spectrum between Batesian and Müllerian mimicry. *Evolution* **31**, 454–455.

Bobisud, L. E. & Potratz, C. J. 1976 One-trial versus multi-trial learning for a predator encountering a model-mimic system. *Am. Nat.* **110**, 121–128.

Bouton, M. E. 1994 Conditioning, remembering and forgetting. *J. Exp. Psychol. Animal Behaviour Processes* **20**, 219–231.

Brower, J. V. Z. 1960 Experimental studies of mimicry. Part IV. The reactions of starlings to different proportions of models and mimics. *Am. Nat.* **94**, 271–282.

Brown, K. S. & Benson, W. W. 1974 Adaptive polymorphism associated with multiple Müllerian mimicry in *Heliconius numata* (Lepid. Nymph.). *Biotropica* **6**, 205–228.

Bush, R. R. & Mosteller, F. 1955 *Stochastic models for learning*. New York: Wiley.

Chai, P. 1986 Field observations and feeding experiments on the responses of rufous-tailed jacamars (*Galbula ruficauda*) to free-flying butterflies in a tropical rain-forest. *Biol. J. Linn. Soc.* **29**, 161–189.

Estabrook, G. F. & Jespersen, D. C. 1974 Strategy for a predator encountering a model-mimic system. *Am. Nat.* **108**, 443–457.

Ford, E. B. 1975 *Ecological genetics*, 4th edn. London: Chapman and Hall.

Gilbert, L. E. 1983 Coevolution and mimicry. In *Coevolution*, ed. D. J. Futuyma and M. Slatkin, pp. 263–281. Sunderland, Massachusetts: Sinauer.

Gleitman, H. 1971 Forgetting of long term memories in animals. In *Animal memory*, ed. W. K. Honig and P. H. R. James, pp. 1–44. London: Academic Press.

Huheey, J. E. 1964 Studies of warning coloration and mimicry. IV. A mathematical model of model-mimic frequencies. *Ecology* **45**, 185–188.

Huheey, J. E. 1976 Studies on warning coloration and mimicry. VII. Evolutionary consequences of a Batesian-Müllerian spectrum; a model for Müllerian mimicry. *Evolution* **30**, 86–93.

Huheey, J. E. 1980*a* Batesian and Müllerian mimicry: semantic and substantive differences of opinion. *Evolution* **34**, 1212–1215.

Huheey, J. E. 1980*b* Studies in warning coloration and mimicry. VIII. Further evidence for a frequency-dependent model of predation. *J. Herpetol.* **14**, 223–230.

Huheey, J. E. 1988 Mathematical models of mimicry. *Am. Nat.* **131**, S22-S41.

Kimler, W. C. 1983*a* One hundred years of mimicry: history of an evolutionary exemplar. Ph.D. thesis, Cornell University, Ithaca, New York.

Kimler, W. C. 1983*b* Mimicry: views of naturalists and ecologists before the Modern Synthesis. In *Dimensions of Darwinism*, ed. M. Grene, pp. 97–127. New York: Cambridge University Press.

Luedeman, J. K., McMorris, F. R. & Warner, D. D. 1981 Predators encountering a model-mimic system with alternative prey. *Am. Nat.* **117**, 1040–1048.

Mackintosh, N. J. 1975 A theory of attention: variations in the associability of stimuli with reinforcement. *Psychol. Rev.* **82**, 276–298.

Mackintosh, N. J. 1983 *Conditioning and associative learning*. Oxford University Press.

MacNamara, J. & Houston, A. I. 1987 Memory and the efficient use of information. *J. theor. Biol.* **125**, 385–395.

Owen, R. E. & Owen, A. R. G. 1984 Mathematical paradigms for mimicry: recurrent sampling. *J. theor. Biol.* **109**, 217–247.

Pasteur, G. 1982 A classificatory review of mimicry systems. *A. Rev. Ecol. Syst.* **13**, 169–199.

Pavlov, I. P. 1927 *Conditioned reflexes*. New York: Oxford University Press.

Pearce, J. M. & Hall, G. 1980 A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* **87**, 532–552.

Rescorla, R. A. & Wagner, A. R. 1972 A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning*, vol. 2 (*Current research and theory*), ed. A. H. Black & W. F. Prokasy, pp. 64–99. New York: Appleton-Century-Crofts.

Rothschild, M. 1980 Mimicry, butterflies and plants. *Symb. bot. Upsal.* **22** (4), 82–89.

Rothschild, M. 1981 The mimicrats must move with the times. *Biol. J. Linn. Soc.* **16**, 21–23.

Sbordoni, V., Bullini, L., Scarpelli, G., Forestiero, S. & Rampini, M. 1979 Mimicry in the burnet moth *Zygaena ephialtes*: population studies and evidence of a Batesian-Müllerian situation. *Ecol. Entomol.* **4**, 83–93.

Sheppard, P. M. 1975 *Natural selection and heredity*, 4th edn. London: Hutchinson.

Sheppard, P. M. & Turner, J. R. G. 1977 The existence of Müllerian mimicry. *Evolution* **31**, 452–453.

Speed, M. P. 1990 Mimicry and the psychology of predation. Ph.D. thesis, University of Leeds.

Speed, M. P. 1993 Muellerian mimicry and the psychology of predation. *Anim. Behav.* **45**, 571–580.

Speed, M. P. & Turner, J. R. G. 1996 Learning and forgetting in mimicry. II. Simulations of the natural mimicry spectrum. *Phil. Trans. R. Soc. Lond.* B (Submitted).

Turner, J. R. G. 1977 Butterfly mimicry: the genetical evolution of an adaptation. *Evol. Biol.* **10**, 163–206.

Turner, J. R. G. 1984*a* Darwin's coffin and Doctor Pangloss – do adaptationist models explain mimicry? In *Evolutionary ecology*, ed. B. Shorrocks (British Ecological Society Symposium No. 23), pp. 313–361. Oxford: Blackwell Scientific Publications.

Turner, J. R. G. 1984*b* Mimicry: the palatability spectrum and its consequences. In *The biology of butterflies*, ed. R. I. Vane-Wright & P. R. Ackery (Royal Entomological Society of London Symposium No 11), pp. 141–161. London: Academic Press.

Turner, J. R. G. 1987 The evolutionary dynamics of batesian and muellerian mimicry: similarities and differences. *Ecol. Entomol.* **12**, 81–95.

Turner, J. R. G. 1995 Mimicry as a model for coevolution. In *Biodiversity and evolution*, ed. R. Arai, M. Kato & Y. Doi, pp. 131–150. Tokyo: National Science Museum Foundation.

Turner, J. R. G., Kearney, E. P. & Exton, L. S. 1984 Mimicry and the Monte Carlo predator: the palatability spectrum and the origins of mimicry. *Biol. J. Linn. Soc.* **23**, 247–268.

Vane-Wright, R. I. 1976 A unified classification of mimetic resemblances. *Biol. J. Linn. Soc.* **8**, 25–56.

Wagner, A. R. 1981 SOP: A model of automatic memory processing in animal behaviour. In *Information processing in animals: conditioned inhibition*, ed. N. E. Spear & R. R. Miller, pp. 5–47. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Wagner, A. R. & Rescorla, R. A. 1972 Inhibition in Pavlovian conditioning: application of a theory. In *Inhibition and learning*, ed. R. A. Boakes & M. S. Halliday, pp. 301–336. London: Academic Press.

Wickler, W. 1968 *Mimicry in plants and animals*. London: Weidenfeld and Nicolson.

Received 17 August 1995; accepted 7 December 1995