PROCEEDINGS OF THE ROYAL SOCIETY | **B BIOLOGICAL SCIENCES**

Components of change in the evolution of learning and unlearned preference

Aimee S. Dunlap and David W. Stephens

References	This article cites 18 articles, 2 of which can be accessed free http://rspb.royalsocietypublishing.org/content/early/2009/06/12/rspb.2009.0602.full.html#ref-list-1
P<P	Published online 17 June 2009 in advance of the print journal.
Subject collections	Articles on similar topics can be found in the following collections behaviour (547 articles) cognition (146 articles) evolution (768 articles)
Email alerting service	Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click **here**

Advance online articles have been peer reviewed and accepted for publication but have not yet appeared in the paper journal (edited, typeset versions may be posted when available prior to final publication). Advance online articles are citable and establish publication priority; they are indexed by PubMed from initial publication. Citations to Advance online articles must include the digital object identifier (DOIs) and date of initial publication.

To subscribe to *Proc. R. Soc. B* go to: **http://rspb.royalsocietypublishing.org/subscriptions**

Components of change in the evolution of learning and unlearned preference

Aimee S. Dunlap* and David W. Stephens

Department of Ecology, Evolution, and Behavior, University of Minnesota—Twin Cities, Saint Paul, MN, USA

Several phenomena in animal learning seem to call for evolutionary explanations, such as patterns of what animals learn and do not learn. While several models consider how evolution should influence learning, we have very little data testing these models. Theorists agree that environmental change is a central factor in the evolution of learning. We describe a mathematical model and an experiment, testing two components of change: reliability of experience and predictability of the best action. Using replicate populations of *Drosophila* we varied statistical patterns of change across 30 generations. Our results provide the first experimental demonstration that some types of environmental change favour learning while others select against it, giving the first experimental support for a more nuanced interpretation of the selective factors influencing the evolution of learning.

Keywords: evolution; learning; Drosophila; experimental evolution

1. INTRODUCTION

Learning is a fundamental behavioural process. Recognizing this, investigators ranging from psychologists to molecular neurobiologists have studied learning with impressive results. These studies describe the phenomena and underlying mechanisms of learning. In contrast to this rich body of work, we have few hard facts about how learning evolves. Although we have a handful of empirical papers that illustrate the fitness consequences of learning (e.g. Sullivan 1988; Hollis *et al.* 1997; Mahometa & Domjan 2005), and a compelling body of comparative work on species and sex differences (e.g. Balda *et al.* 1996; Lefebvre 1996; Dunlap *et al.* 2006), models and speculations dominate thinking about the central question of how learning evolved. This is in part because the key explanatory variables of these models have seemed experimentally inaccessible.

To understand the variables that have seemed so problematic, consider why learning may make more sense in some situations than others. Although many models have addressed this question, they all focus on the nature of statistical relationships between stimuli and consequences in the animal's environment (e.g. Stephens 1991, 1993; Papaj 1994; Kerr & Feldman 2003). The simplest and oldest of these ideas focuses on change. Learning, the argument goes, exists because environments change and it follows that animals must use experience to adjust to this change (e.g. Thorpe 1963; Plotkin & Odling-Smee 1979; Johnston 1982; Papaj & Prokopy 1989). This logic leads to a simple claim that we call the 'learning folk theorem': changing environments favour learning, but stable fixed environments favour non-learning (innate or fixed behaviour) (e.g. Plotkin & Odling-Smee 1979; Anderson 1995; Mery & Kawecki 2004).

While the folk theorem continues to influence the thinking of casual students of learning evolution, recent models argue that it oversimplifies the problem (Stephens 1991; Kerr & Feldman 2003; Borenstein *et al.* 2008).

According to these models, the folk theorem mistakenly lumps all forms of environmental change together, when in reality—these models argue—some components of environmental change, such as between-generational change, select for learning (as the folk theorem suggests), but others select against learning. Regardless of which model one favours, controlling or observing the statistical relationships in an animal's environment presents a significant empirical challenge. This paper develops a simple, experimentally tractable 'components of change' model, and presents results from a study testing this model experimentally. This experiment varies relevant components of 'environmental change' across many generations and it offers the first experimental confirmation of the claim that some types of change promote learning while others promote non-learning.

To develop our model and test its predictions, we focus on an experimental preparation developed by Mery & Kawecki (2002, 2004). In this preparation, the experimenter presents two types of egg-laying media to a small group of female *Drosophila melanogaster*: one option is flavoured with orange juice and the other is flavoured with pineapple. The experimenter offers this choice twice. In the first presentation, the investigator pairs one of the media types with the aversive chemical quinine, so the flies experience either (i) orange plus quinine versus pineapple without quinine, or (ii) orange without quinine versus pineapple plus quinine. We call this first presentation the 'experience' phase, because flies experience the pairing with quinine at this stage. In the second presentation, the experimenter again offers the orange and pineapple media, but now without quinine in either media type. We call this second presentation the consequence phase, because—as we will explain—this is when the investigator imposes a fitness consequence. Note that this arrangement creates a relatively simple aversion learning problem in which a 'learning' fly can use the experience of quinine pairing in the first phase (hence the name experience phase) to adjust its egg-laying behaviour in the second phase.

*Author for correspondence (dun10063@umn.edu).

1

Using this preparation, an experimenter can control two variables that, according to our components of change model, influence the fitness value of aversion learning: the reliability of experience and the fixity of the best action. First, the experimenter creates the next generation of flies by rearing eggs from one of the media types in the consequence phase (eggs laid in the experience phase are discarded). The investigator can, for example, create an 'orange-best' situation by rearing only those eggs laid on the orange media. Second, the experimenter can control the extent to which the quinine pairing in the experience phase reliably indicates the best action in the consequence phase.

Mery and Kawecki used this preparation to test the learning folk theorem. They created a changing environment that should favour learning by alternating orange best (rear eggs only from orange) and pineapple best from one generation to the next. In addition, pairing with quinine in the experience phase reliably indicated the media type that flies should avoid in the consequence phase. In agreement with the folk theorem, Mery and Kawecki found enhanced learning (that is, increased sensitivity to the experience of quinine pairing) in this changing environment. In a second study, Mery and Kawecki created a fixed environment in which they always reared eggs (for example) from the orange media type in the consequence phase. Contrary to their expectations, Mery and Kawecki also found increased responsiveness to experience in this condition. Here, again, pairing with quinine in the experience phase reliably indicated the best action in the consequence phase. As we will explain in the following, according to the components of change view of learning, Mery and Kawecki's 'fixed environment' situation does not, in fact, select against learning. This is because while fixity of the best action does select against learning as the folk theorem claims, fixity of the relationship between experience and consequence favours learning.

2. THE COMPONENTS OF CHANGE MODEL

(a) *The parameters p and q*

Here, we develop a model based on the Mery–Kawecki preparation (the appendix presents the algebraic details). Let p represent the overall probability that the experimenter rears eggs from the orange flavoured media (so that laying eggs on orange is the best action). We focus on 'orange' to simplify the model development; focusing on pineapple yields identical results. The parameter p ($0.5 \leq p \leq 1.0$) specifies the fixity of the best action, and we call it the *best-action fixity*. This is our first component of change. For example, $p = 1.0$ gives the highest possible best-action fixity because it means that the experimenter always rears eggs from orange, and 'lay eggs on orange' is always the best policy. In contrast, $p = 0.5$ gives the lowest meaningful value of best-action fixity because it means that laying eggs on orange is the best half the time and laying on pineapple is the best half the time. Similarly, we use the parameter q to represent the fixity of the relationship between experience and best action. Mathematically, q is the conditional probability that the experimenter rears eggs from the substrate type that was NOT paired with quinine in the first or experience phase of the experiment. The parameter q ($0.5 \leq q \leq 1.0$)

therefore measures the fixity of the relationship between experience and the best action. As before we simplify the terminology by calling this variable the *reliability of experience*, which is our second component of change. If $q = 1$, the flies can reliably select the best action by avoiding the substrate that was paired with quinine in the experience phase; however if $q = 0.5$ pairing with quinine carries no information about the fitness consequences of egg-laying choices in the second (or consequence) phase of the experiment.

(b) *Model predictions: learning versus non-learning*

To evaluate the effects of these parameters we compare the fitness of a non-learning genotype that always lays eggs on orange (because we have arbitrarily assumed that 'lay on orange' is the most common best action, i.e. $0.5 < p < 1.0$) to the fitness of a learning genotype that uses the pairing with quinine in the experience phase to guide its behaviour in the 'consequence' phase. Figure 1 shows the results of these calculations. The figure shows how the two fixity parameters influence the relative fitness obtained by our learning and non-learning genotypes. As the figure shows, a diagonal line (running from $(p, q) = (0.5, 0.5)$ to $(1, 1)$) separates the learning and non-learning regions; the learning genotype does better above the line while the non-learning genotype should prevail below the line. For example, the point where $p = 0.5$ and $q = 1.0$ strongly promotes learning because the best action changes randomly (i.e. there is low best-action fixity, $p = 0.5$), yet the experience of quinine pairing reliably signals the best action (i.e. there is high reliability of experience, $q = 1.0$). This crudely corresponds to Mery & Kawecki's (2002) 'changing environment' study, in which they found that learning abilities improved within 20 generations of selection. Notice however that the condition that most strongly selects against learning is the point where there is high best-action fixity ($p = 1.0$) and low reliability of experience ($q = 0.5$). Mery and Kawecki did not test this situation. Instead they tested the 'completely fixed' case (i.e. $p = q = 1.0$), which, following the 'folk theorem' they argued, should select against learning. However, as the figure shows, this situation is actually selectively neutral (see the appendix for mathematical rationale). In the absence of a learning cost, complete fixity neither favours nor disfavours our learning genotype.

3. EXPERIMENTAL APPROACH AND METHODS

The experiment presented here compares selection in the two regimes that most strongly favour and disfavour learning. We assigned small populations of *D. melanogaster* to three conditions. (i) High best-action fixity ($p = 1.0$) and low reliability of experience ($q = 0.5$): this strongly disfavours learning because experience is unreliable and the same action is always best. (ii) Low best-action fixity ($p = 0.5$) and high reliability of experience ($q = 1.0$): this strongly favours learning because sensitivity to the experience of quinine pairing allows flies to consistently track the best action. (iii) In addition, we established control populations with the same initial population size and rearing procedures as the two experimental groups, but these flies never experienced the fruit-flavoured media
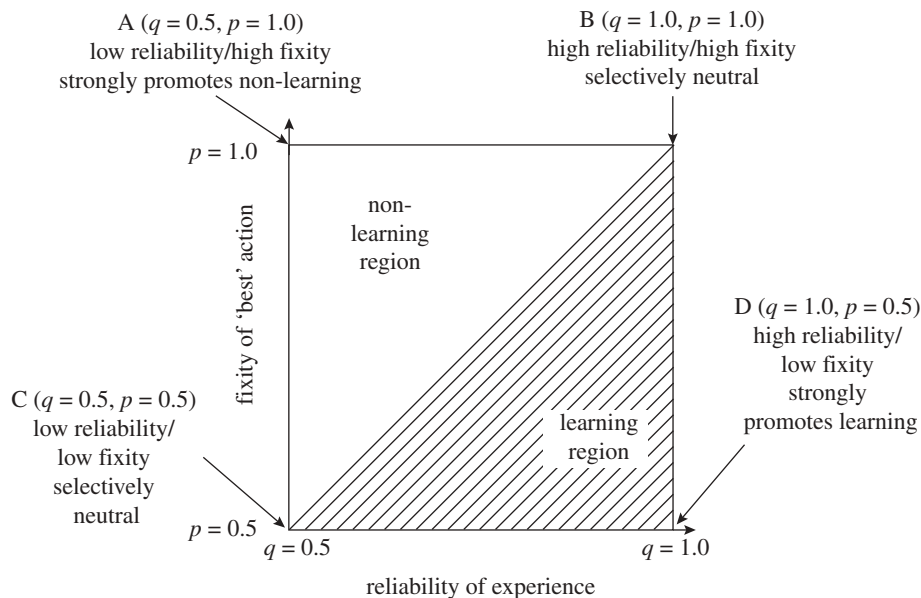
Figure 1. Best-action fixity and reliability of experience influence the fitness value of learning. Natural selection favours learning most strongly when the reliability of experience is high, but best-action fixity is low (Point D); selection favours non-learning when best-action fixity is high, but reliability is low (Point A). We test points A and D.

or quinine. Note that one can, in principle, fix the best action in two ways: orange always best or pineapple always best. We included both possibilities in our design by randomly assigning half of the populations in each treatment to 'orange initially' best and half to 'pineapple initially' best conditions. Although we made this assignment for all three treatments, it has different implications for each of the three treatments. For the 'high best-action fixity' treatment, it fixes the best action (lay on orange or lay on pineapple) across all 30 generations; for the 'low best-action fixity', it determines the initial state, but this changes randomly in subsequent generations; for the control lines—which never experienced orange or pineapple—it is simply an arbitrary designation.

(a) Initial populations and treatment assignments

To create our initial stock of flies, we mixed 400 males and 400 females from each of four lab-adapted, wild-caught populations from Minnesota and Wisconsin (USA). We maintained them in overlapping generations in a large population cage for five months prior to the start of the experiment. We housed all flies at 24°C. We reared all eggs at a density of 80 eggs per vial, with six vials per line per generation. We established 36 lines of 400 flies and randomly assigned 12 lines to each of three treatments. For each treatment, we randomly assigned 6 of the 12 lines to orange best and 6 to 'pineapple best'. As outlined earlier, our three selection treatments are: (i) 'best-action fixed ($p = 1.0$)/experience unreliable ($q = 0.5$)'; (ii) 'best-action changing ($p = 0.5$)/experience reliable ($q = 1.0$)' treatment; and (iii) control.

(b) Implementing the experimental selection regimes

Every generation, we transferred 200 female flies (along with a similar number of males) from each line to a test cage. The cages were approximately shoe-box size (33.3 cm length × 21 cm width × 12 cm height), and we

equipped each of them with a sliding drawer that could hold two petri dishes. We presented petri dishes with standard cornmeal and molasses media until we were ready to begin the selection (3 days). As the introduction explains, the experimental selection regime consists of two phases: an experience phase (in which we paired quinine with one type of media) and a consequence phase (where quinine was never present). The experience phase exposed flies to two fruit flavours of agar-based media in a single 3 h session (reconstituted frozen orange or pineapple juice, 12 g agar/1 l juice, with 20 ml of juice agar placed in the bottom of each 100 mm × 15 mm petri dish). Following our experimentally determined schedule, we paired quinine with one of the two flavours (4 g quinine/1 l agar). In the consequence phase, we presented fresh petri dishes of the two flavours of media (using the sliding drawer to change the media). We randomized the locations of orange and pineapple plates within each cage, but kept the location the same in the experience and consequence phases for a given line in a given generation. An interval of 30 min separated the removal of the experience phase plates and the introduction of the consequence phase plates.

Following an experimentally determined schedule, we reared eggs laid on only one of the media flavours in the consequence phase, and discarded all other eggs. We removed eggs selected for propagation from the substrate using a needle and placed them in vials on standard cornmeal-based fly food for incubation.

(c) End of selection assays

Following 30 generations of selections, we tested each line in a series of assays. We reared the flies used in these assays from eggs collected on standard (unflavoured) media. We conducted two types of assays: learning assays and preference (or non-learning) assays. Our learning assays consisted of two tests. First, we tested a group of 200 naive females (with a similar number of males) from each line with a 3 h experience phase of quinine

paired with orange. Second, we tested a different group of naive flies with quinine paired with pineapple. We followed both with a 1.5 h consequence phase in which neither flavour was paired with quinine. (Note that we use the terminology 'experience phase' and 'consequence phase' for simplicity here, even though these assays differ in some details from the experimental selection procedures.) Our preference assay tested flies with no quinine present during a 3 h experience phase, and no quinine present during the 1.5 h consequence phase.

4. RESULTS

(a) *Evolution of non-learning (preference)*

We tested the effect of our three selection regimes on unlearned preferences by simultaneously presenting orange and pineapple flavoured media to naive groups of flies from each of our treatments, and observing the number of eggs laid on each type of media. Figure 2 shows these data expressed as the proportional preference for the orange media. As our model predicts, the figure shows changes in unlearned preferences for the best-action fixed ($p = 1.0$)/experience unreliable ($q = 0.5$) lines, but not for the best-action changing ($p = 0.5$)/experience reliable ($q = 1.0$) treatment. Focusing on the best-action fixed/experience unreliable lines, we see a striking difference between the lines assigned to the orange-best and pineapple-best conditions. As we would expect, lines assigned to the pineapple-best treatments showed a decreased preference for orange. An analysis of variance of all three selection regimes supports this interpretation by showing a significant interaction between treatment and best assignment ($F_{2,30} = 3.381$, $P = 0.0474$). In addition, post hoc analyses show a difference between the orange-best and pineapple-best lines for the best-action fixed/experience unreliable treatments but not for the other treatments.

(b) *Evolution of learning*

To assess differences in sensitivity to experience, we exposed groups of flies from each of our treatments to an assay that closely paralleled our selection procedures. In this assay, we paired quinine with either orange or pineapple, and then scored oviposition preferences in a second presentation of the two types of media without quinine. By testing a separate group of flies from each line in both an orange paired with quinine and a pineapple paired with quinine condition, we can derive a contingency score for each line using Cramer's ϕ. This score measures the extent to which oviposition preferences in the second stage of the assay depend on the experience of quinine pairing in the first stage. Figure 3 shows these data. As our model predicts, we see enhanced sensitivity to experience in the best-action changing ($p = 0.5$)/experience reliable ($q = 1.0$) treatment (compared with the control group) and no difference in sensitivity to experience between the best action fixed/experience unreliable treatment and the control. A one-way analysis of variance confirms a significant effect of treatment ($F_{2,33} = 4.17$, $P = 0.02$). In addition, post hoc analyses (using Tukey's LSD) confirm the pattern shown in the figure. Specifically, the best-action changing ($p = 0.5$)/experience reliable ($q = 1.0$) treatment shows a higher sensitivity to experience (as measured by ϕ) than either the control or best-action fixed ($p = 1.0$)/experience

unreliable ($q = 0.5$) treatments. We find the same statistically significant results in a more complex analysis of the proportion of eggs laid on the substrate consistent with learning (best-action changing/experience reliable, $\overline{X} \pm$ SD 0.698 ± 0.099; best-action fixed/experience unreliable, $\overline{X} \pm$ SD 0.635 ± 0.162; control $\overline{X} \pm$ SD 0.645 ± 0.150).

(c) *General selection results*

Flies never oviposited on a substrate when quinine was present, and this avoidance of quinine was the same for all treatments, and it did not change during the course of selection. Although less tidy, an analysis of data from selection trials is in broad agreement with the analyses presented above. Specifically, we calculated the extent to which flies avoided the media type that had been paired with quinine in the experience phase of selection trials using the proportion of all eggs laid on this type of media (a variable we call P(response to experience)). To account for changes in these measures across generations, we calculated P(response to experience) values for each line in blocks of two generations each. Finally, we analysed these scores in an ANOVA with factors of treatment and blocks, with repeated measures on each line. This analysis showed main effects of treatment ($F_{1,22} = 6.51$, $P = 0.018$), with best-action changing ($p = 0.5$)/experience reliable ($q = 1.0$) lines showing higher learning scores than best-action fixed ($p = 1.0$)/experience unreliable ($q = 0.5$) lines; and a main effect of block is also statistically significant ($F_{4,308} = 2.31$, $P = 0.005$), but the interaction between the two is not quite significant ($F_{14,308} = 1.62$, $P = 0.071$).

5. DISCUSSION

This study offers an experimental analysis of the selective value of learning. Specifically, it asks how two components of change (the reliability of experience, and underlying uncertainty about the appropriate action) affect the value of learning. It is, to our knowledge, the first experimental confirmation of the insight that that these two statistical relationships can select both for and against learning. Our result illustrates the weakness of the influential claim of the learning folk theorem that 'change favours learning' while 'fixity favours non-learning'. Our results suggest that randomness, and not fixity, is the most powerful and plausible way to select against learning. Consider, for example, the Garcia effect (Garcia & Koelling 1966), which shows that rats learn associations between tastes and gastric illness more easily than the association between bright–noisy and gastric illness. Surely this does not happen because the relationship between visual stimuli and gastric illness has been fixed throughout rat evolutionary history. It is much more plausible that visual stimuli have varied unpredictably in relation to gastric consequences.

Our study, of course, owes much to the ground-breaking work of Mery & Kawecki (2002, 2004). Mery and Kawecki's two studies using this experimental system tested the role of change in the evolution of learning, motivated by the learning folk theorem. Our best-action changing/experience reliable treatment replicates Mery and Kawecki's first study in that both studies found that this condition selected for enhanced learning. Our study, however, introduced random change while Mery
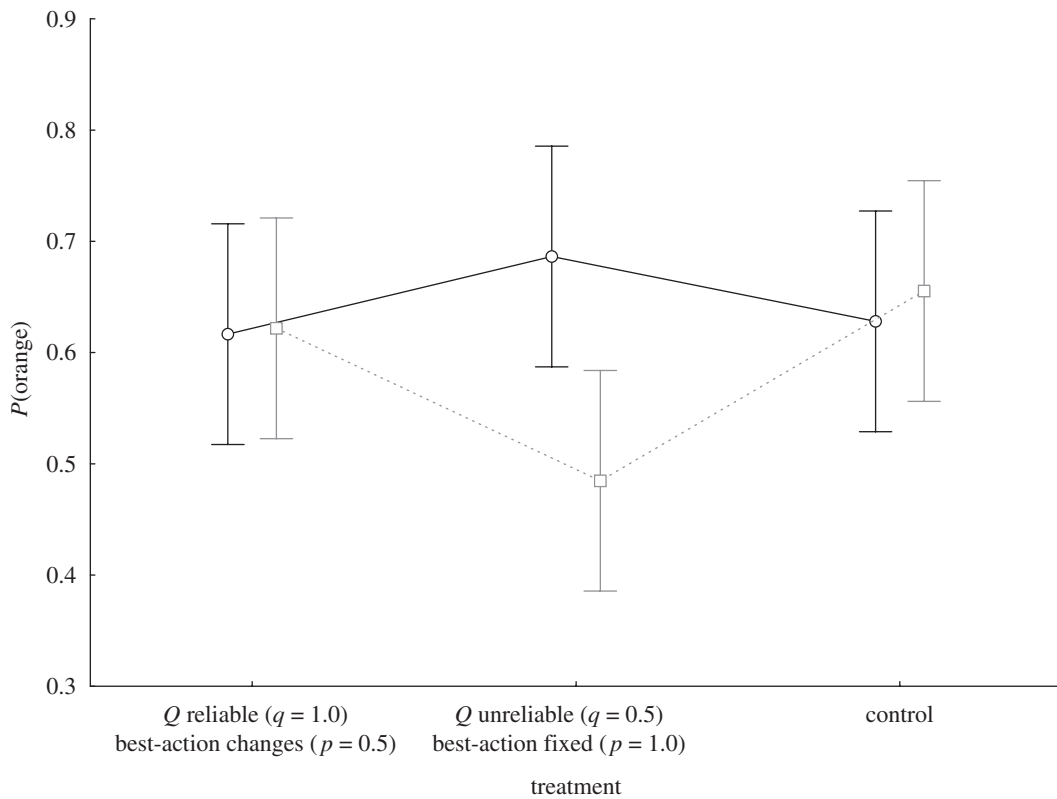
Figure 2. Interaction of treatment by initial best assigned environment during following the 30th generation of selection. P(orange) is the proportion of eggs laid on orange during a test of preference without quinine experience. Error bars are 95 per cent confidence intervals. Solid line, open circle, orange initially best; dotted line, open square, pineapple initially best.
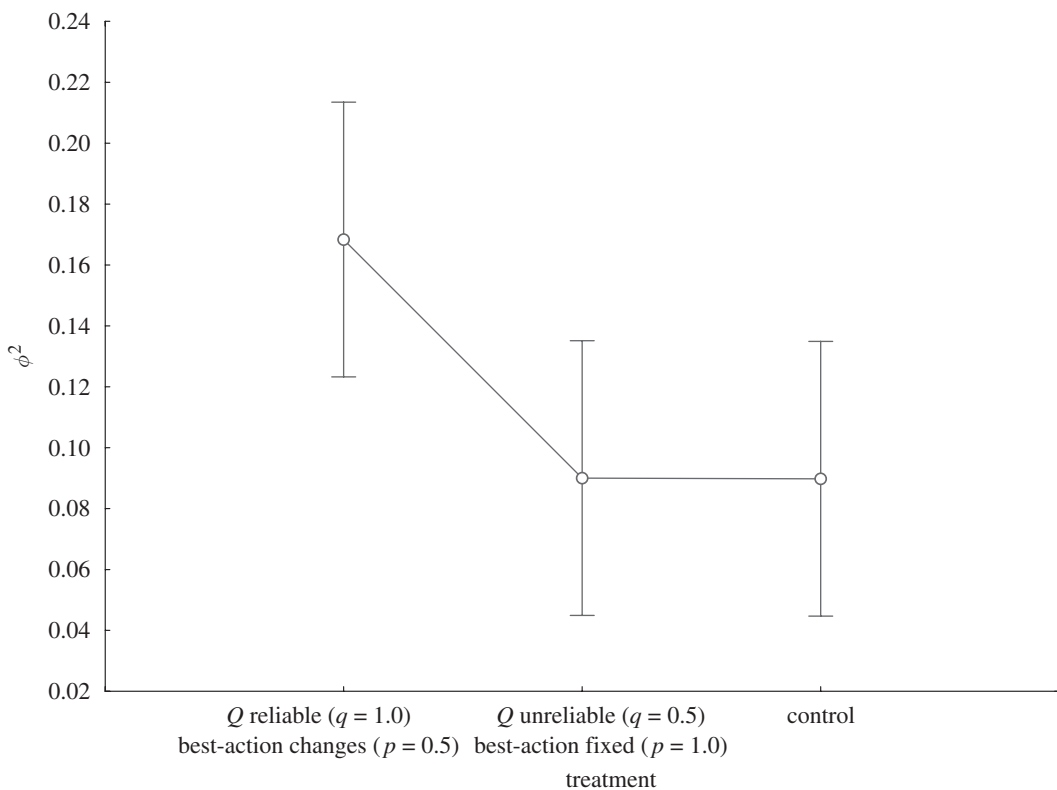

Figure 3. Significant effect of treatment during the final experiment assays following the 30th generation of selection. ϕ^2 is a measure of contingency of the effect of quinine experience. Error bars are 95 per cent confidence intervals. The difference between reliable/unpredictable and the other treatments is statistically significant.

and Kawecki strictly alternated orange-best and pineapple-best conditions. And although we expect differences in generalized learning abilities, as Mery and Kawecki found, we did not explicitly test this. The key difference between our approaches, however, follows from different perspectives about the condition that selects against learning. Following the learning folk theorem, Mery and Kawecki tested an absolutely fixed condition in which the best action was always the same, and where quinine reliably predicted the best action. Contrary to their expectation, they found enhanced learning in this situation. In contrast, following the components of change view of learning, we tested a condition where the best action was always the same, but where there was no predictable relationship between quinine pairing and the best action. As predicted, we find reduced sensitivity to experience and increased reliance on unlearned preferences in this selection regime.

While Mery and Kawecki's work represents the only similar empirical studies, our work has deep connections to theoretical work on the selective value of learning (Johnston 1982; Stephens 1991; Papaj 1994; Dukas 1998; Kerr & Feldman 2003). As a group, these papers emphasize the role of change and other statistical properties of the environment in learning evolution. Early work by Johnston emphasizes the learning folk theorem, even though it acknowledges that animals should not learn in some changing environments (e.g. under complete unpredictability). The later papers take an increasingly nuanced view that either recognizes different components of change (e.g. Stephens 1991; Papaj 1994) or argues that intermediate levels of change favour learning (Kerr & Feldman 2003). This paper, perhaps unsurprisingly, is most clearly connected to the Stephens (1991) model. The parallels between the Stephens model and the Mery–Kawecki experimental preparation (used here) are striking. Stephens modelled a hypothetical organism with a two-stage life history. In the first stage, the organism can choose to obtain experience; while in the second stage the animal can act in response to its experience in the first stage. Although the Stephens model characterized the components of environmental change in a different way, its predictions closely follow the model presented here with one key difference: the Stephens model predicts non-learning for the absolutely fixed condition. This difference occurs because the Stephens model imposed an opportunity cost on learning. Specifically, in the Stephens model a learner can waste time acquiring experience in the experience phase of its life history, when the analogous non-learner can begin to acquire fitness benefits in the experience phase. This cannot happen in the Mery–Kawecki preparation, because choices made in the experience phase do not affect fitness. Natural learning surely imposes some costs (both opportunity costs and physiological costs), however models suggest—in agreement with our experimental results—that unpredictability is a much more powerful and robust way to select against learning than fixity, even when learning imposes costs.

6. SUMMARY

The experimental analysis presented here exploits Mery and Kawecki's pioneering empirical paradigm to test a logically coherent model of learning evolution. This model recognizes two distinct types of 'fixity' that have opposing effects on the selective value of learning: (i) fixity of the best actions (e.g. it is always best to lay eggs on orange) selects against learning (as the folk theorem claims), and (ii) fixity of the relationship between stimuli and best action (e.g. quinine is always paired with the worst type of media) favours learning. Our results support this more complicated claim. In treatments with a fixed 'best action' and an unreliable (changing) relationship between stimuli and best action, we observed increased non-learning (i.e. simple preferences for media type). On the other hand, in treatments where the best action changed and we created a reliable (fixed) relationship between stimuli and best action, we observed increased learning.

Learning is a fundamental mechanism for adjusting behaviour to change in the environment. Our results emphasize a richer and more realistic view of the evolutionary advantages of this flexibility, recognizing that different components of environmental change can have different effects on the evolution of learning and phenotypic plasticity. This perspective is significant because it is immediately relevant to the explanation of variation in animal learning abilities such the Garcia effect and other examples of selective association in animal learning.

APPENDIX A

To begin, we introduce some notation that simplifies our presentation. First, we use A and B to represent the two types of media (orange and pineapple in our experiment). Second, we use the notation QA to mean that quinine is paired with media type A in the experience phase; similarly QB means that we paired quinine with media type B in the consequence phase. Third, we use the notation A* to mean that the experimenter rears eggs from media type A in the consequence phase; similarly we use B* to mean that we rear eggs from media type B. To simplify the terminology, we say that the A* is the 'A-best' condition; similarly B* is the 'B-best' condition.

Next, we use this notation to define parameters that represent the fixity of the stimulus–consequence and action–consequence relationships. Let q measure the fixity of the stimulus–consequence relationship, specifically $P(A^*|QB) = P(B^*|QA) = q$. In words, q is the conditional probability that pairing with quinine in the experience phase predicts the media that flies should avoid in the consequence phase.

We use p to represent the fixity of action–consequence relationship. Specifically, let p be the probability that the A-best condition applies for any given realization of the consequence phase. For example, if $p = 1.0$, it is

event			non-learner		learner	
experience phase	consequence phase	probability	behaviour	fitness	behaviour	fitness
QA	A*	$p(1-q)$	prefer A	$r(1-\epsilon)n$	prefer B	$r\epsilon n$
QA	B*	$(1-p)q$	prefer A	$r\epsilon n$	prefer B	$r(1-\epsilon)n$
QB	A*	pq	prefer A	$r(1-\epsilon)n$	prefer A	$r(1-\epsilon)n$
QB	B*	$(1-p)(1-q)$	prefer A	$r\epsilon n$	prefer A	$r\epsilon n$

always best to lay eggs of media type A, whereas if $p = 0.5$, the best place to lay varied unpredictably from one generation to the next. We remark that we lose no generality by defining p in terms of the A-best condition, because types A and B are arbitrary. In practice, this means that we define type A to be the type that is 'best' most frequently, i.e. $P(A^*) \geq P(B^*)$, implying that $P(A^*) \geq 0.5$.

Now we consider two types of flies: a non-learner who always tries to oviposit on A and a learner who oviposits on A if quinine was paired with B in the experience phase, but oviposits on B if quinine was paired with A in the experience phase. We assume that a female lays n eggs in the consequence phase. In addition, we assume that a female makes some oviposition errors so that she cannot lay 100 per cent of her eggs in her preferred media. Instead she lays $1 - \epsilon$ of her eggs in the media she 'prefers' and ϵ in the media she 'intends' to avoid; so ϵ is the error rate. Finally, we assume that a proportion r of the eggs a female lays in the 'best' media survive to reproduce, while none survive to reproduce when they are laid on the 'worst' media. Within a generation, there are four possible events, as shown in the table above.

From this, we can calculate the fitnesses of the two types. When fitness varies temporally (from one generation to the next), we calculate the geometric mean fitness of the two types (Karlin & Lieberman 1974, 1975). The fitness of the non-learner is

$$p(1-q)\ln[r(1-\epsilon)n] + (1-p)q\ln(r\epsilon n) + pq\ln[r(1-\epsilon)n]$$
$$+ (1-p)(1-q)\ln(r\epsilon n),$$

which simplifies to

$$P\ln[r(1-\epsilon)n] + (1-p)\ln(r\epsilon n).$$

The reliability term, q, cancels out because the non-learner ignores the pairing with quinine. Similarly, the fitness of the learner is

$$p(1-q)\ln(r\epsilon n) + (1-p)q\ln[r(1-\epsilon)n] + pq\ln[r(1-\epsilon)n]$$
$$+ (1-p)(1-q)\ln(r\epsilon n),$$

which simplifies to

$$Q\ln[r(1-\epsilon)n] + (1-q)\ln(r\epsilon n).$$

Here, the frequency of the A-best state cancels out because the learner's fitness depends on whether the quinine cue reliably predicts the best media. The only difference between the two simplified expressions is the presence of p or of q, thus learning should be favoured

whenever $q > p$. The learning and non-learning traits will be neutral whenever $q = p$. This includes the so-called absolute fixity case—$q = p = 1$—and the completely random case—$q = p = 1/2$.

REFERENCES

Anderson, R. W. 1995 Learning and evolution: a quantitative genetics approach. *J. Theor. Biol.* **175**, 89–101. (doi:10.1006/jtbi.1995.0123)

Balda, R. P., Kamil, A. C. & Bednekoff, P. A. 1996 *Predicting cognitive capacity from natural history*, vol. 13 (eds V. N. Nolan & E. D. Ketterson), pp. 33–66. New York, NY: Plenum Press.

Borenstein, E., Feldman, M. W. & Aoki, K. 2008 Evolution of learning in fluctuating environments: when selection favors both social and exploratory learning. *Evolution* **62**, 586–602. (doi:10.1111/j.1558-5646.2007.00313.x)

Dukas, R. 1998 Evolutionary ecology of learning. In *Cognitive ecology: the evolutionary ecology if information processing and decision making* (ed. R. Dukas), pp. 129–164. Chicago, IL: University of Chicago Press.

Dunlap, A. S., Chen, B. B., Bednekoff, P. A., Greene, T. G. & Balda, R. P. 2006 A state-dependent sex difference in spatial memory in pinyon jays, *Gynorhinus cyanocephalus*: mated females forget as predicted by natural history. *Anim. Behav.* **72**, 401–411. (doi:10.1016/j.anbehav.2006.01.015)

Garcia, J. & Koelling, R. A. 1966 Relation of cue to consequence in avoidance learning. *Psychon. Sci.* **4**, 123–124.

Hollis, K. L., Pharr, V. L., Dumas, M. J., Britton, G. B. & Field, J. 1997 Classical conditioning provides paternity advantage for territorial male blue gouramis (*Trichogaster trichopterus*). *J. Comp. Psychol.* **111**, 219–225. (doi:10.1037/0735-7036.111.3.219)

Johnston, T. D. 1982 Selective costs and benefits in the evolution of learning. *Adv. Study Behav.* **12**, 65–106. (doi:10.1016/S0065-3454(08)60046-7)

Karlin, S. & Lieberman, U. 1974 Random temporal variation in selection intensities: case of large population size. *Theor. Popul. Biol.* **6**, 355–382. (doi:10.1016/0040-5809(74)90016-1)

Karlin, S. & Lieberman, U. 1975 Random temporal variation in selection intensities: one-locus two-allele model. *J. Math. Biol.* **2**, 1–17. (doi:10.1007/BF00276012)

Kerr, B. & Feldman, M. W. 2003 Carving the cognitive niche: optimal learning strategies in homogeneous and heterogeneous environments. *J. Theor. Biol.* **220**, 169–188. (doi:10.1006/jtbi.2003.3146)

Lefebvre, L. 1996 Ecological correlates of social learning: problems and solutions for the comparative method. *Behav. Processes* **35**, 163–171. (doi:10.1016/0376-6357(95)00054-2)

Mahometa, M. J. & Domjan, M. 2005 Classical conditioning increases reproductive success in Japanese quail, *Coturnix*

japonica. Anim. Behav. **69**, 983–989. (doi:10.1016/j. anbehav.2004.06.023)

Mery, F. & Kawecki, T. J. 2002 Experimental evolution of learning ability in fruit flies. *Proc. Natl Acad. Sci. USA* **99**, 14 274–14 279. (doi:10.1073/pnas. 222371199)

Mery, F. & Kawecki, T. J. 2004 The effect of learning in experimental evolution of resource preference in *Drosophila melanogaster. Evolution* **58**, 757–767.

Papaj, D. R. 1994 Optimizing learning and its effect on evolutionary change in behavior. In *Behavioral mechanisms in evolutionary ecology* (ed. L. A. Real), pp. 133–153. Chicago, IL: University of Chicago Press.

Papaj, D. R. & Prokopy, R. J. 1989 Ecological and evolutionary aspects of learning in phytophageous insects. *Annu. Rev. Entomol.* **34**, 315–350. (doi:10.1146/annurev.en.34. 010189.001531)

Plotkin, H. C. & Odling-Smee, F. J. 1979 Learning, change, and evolution: an enquiry into the teleonomy of learning. *Adv. Study Behav.* **10**, 1–41. (doi:10.1016/S0065-3454 (08)60092-3)

Stephens, D. W. 1991 Change, regularity, and value in the evolution of animal learning. *Behav. Ecol.* **2**, 77–89. (doi:10.1093/beheco/2.1.77)

Stephens, D. W. 1993 Learning and behavioral ecology: incomplete information and environmental predictability. In *Insect learning: ecological and evolutionary perspectives* (eds D. R. Papaj & A. Lewis), pp. 195–217. New York, NY: Chapman and Hall.

Sullivan, K. A. 1988 Age-specific profitability and prey choice. *Anim. Behav.* **36**, 613–615. (doi:10.1016/S0003-3472(88)80038-1)

Thorpe, W. H. 1963 *Learning and instinct in animals.* London, UK: Methuen.