Reply

Connecting behavioural biologists and psychologists: Clarifying distinctions and suggestions for further work

John M.C. Hutchinson *, Gerd Gigerenzer

Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

Abstract

This article is a reply to the commentaries on our target article, which relates our group's work on simple heuristics to biological research on rules of thumb. Several commentators contrasted both these approaches with behaviour analysis, in which the patterns of behaviour investigated in the laboratory are claimed to be near-universal attributes, rather than specific to particular appropriate environments. We question this universality. For instance, learning phenomena such Pavlovian or operant conditioning have mostly been studied only in a few generalist species that learn easily; in many natural situations the environment hinders learning as an adaptive strategy. Other supposedly general phenomena such as impulsiveness and matching are outcome models, which several different models of simple cognitive processes might explain. We clarify some confusions about optimisation, optima and optimality modelling. Lastly, we say a little more about how heuristics might be selected, learnt and tuned to suit the current environment.
© 2005 Elsevier B.V. All rights reserved.

Keywords: Animal cognition; Ecological rationality; Impulsiveness; Optimality modelling; Rule of thumb; Simple heuristic

We find lots in the commentaries that will stimulate further research and help establish common ground between our Centre for Adaptive Behaviour and Cognition (ABC) and animal research. For instance, we appreciate Kyonka and Church's ideas about how heuristics like Take The Best could be tested on animals using operant techniques. Another instance is how Shettleworth uses existing results from animal psychology to encourage us in a broadening of perspective from rules of cue usage to rules of cue

learning. Rather than pick out isolated points with which we take issue, we will concentrate on themes that were the concern of several commentaries.

1. Universal versus situation-specific mechanisms

ABC has emphasised that because a heuristic's performance depends on the environment, we expect different heuristics to be used in environments with different statistical structures. Heuristics are thus general only in the sense that the same heuristics might be used in a variety of domains with common statistical structures and by a variety of species. Three commentaries

* Corresponding author. Tel.: +49 30 82406352;
fax: +49 30 82406394.
E-mail address: hutch@mpib-berlin.mpg.de
(J.M.C. Hutchinson).

(Mazur, Williams, Shettleworth) contrast this with the much more general mechanisms that behaviour analysis has emphasised. Mazur makes claims of universality for certain mechanisms of learning, whereas Williams does admit of limited situation-specific exceptions such as bird song.

If these were indeed universal constraints about how brains work, there would be no problem with heuristics being designed around them; cognitive constraints exist whether or not any are universal. However, we would claim that Mazur and Williams underestimate the exceptions. For instance, not all stimuli-response associations are equally readily learnt, sometimes in ways that make adaptive sense: Shettleworth mentions Garcia and Koelling's (1966) work on which cues to toxicity rats can learn. Claims of universality may actually be based on rather few species, and, as Williams mentions, it may be a deliberately biased sample of generalist species that can learn how to use Skinner boxes.

The data are also biased in that learning is inevitably studied in situations where it is easy to learn. In many natural situations individual learning, such as Pavlovian or operant conditioning, is difficult because events are rare, single errors are dangerous, lifetimes are short, or because it is difficult to get immediate feedback about the consequences of a choice. Also complex situations, as Mazur mentions involving "multiple cues and multiple dimensions, and multiple source of uncertainty", make the interpretation of feedback much more difficult than in the simpler well-controlled world of a Skinner box. In some situations, the past may even be an unreliable guide to the future (for instance, with a non-renewing evenly distributed resource, finding an item should be the stimulus *not* to look there again). In all these situations, individual learning need not be as ecologically rational as alternatives, such as a hardwired response (see Cross and Jackson's discussion of "Darwinian" animals) or, especially in humans, a reliance on social heuristics such as doing what the majority of your peers do (Laland, 2001). Thus, whereas Pitts envisages that triage doctors might learn valid cues via a history of differential reinforcement, research actually finds that opportunities for feedback are rare and that doctors instead mostly apply rules taught in medical school (Gigerenzer, 2002). We would therefore claim that the supposedly universal mechanisms of individual learning are applicable only in particular types of environments, although perhaps these are widespread ones. It may help to connect ABC's work with behaviour analysis by considering that each concentrates on a different region along a continuum between mechanisms that are more and less widely applicable. One lesson of such a perspective for students of learning is that they should establish in what environments individual learning is ecologically rational compared, say, with a hardwired response.

2. Outcome versus process models

Amongst the other supposedly universal laws mentioned by Williams and Mazur, impulsiveness and matching contrast with heuristics in another way: they are descriptions of the behaviour (outcome models or as-if models) rather than models of the cognitive processes that generate the behaviour. Sanabria and Killeen briefly make this distinction and our viewpoint is also consistent with Pitts' doubts that matching is itself an evolved heuristic but rather results from an interaction between a particular evolved susceptibility to stimuli and the environment.

As an illustration, Thuijsman et al. (1995) have provided two examples of simple heuristics that would enable a bee to monitor the rewards available from two species of flower and prefer the one producing more nectar. One heuristic is the "ε-Sampling Strategy": stick with a single species except with some low probability ε sample the other species and then switch if the payoff exceeds an average of recent payoffs. The other even simpler heuristic is the "Failures Strategy" in which a bee switches species whenever it has experienced *n* empty flowers in a row. When a bee is studied in isolation both heuristics (and many other learning rules) can produce matching behaviour. Incidentally, although matching in this one-bee environment is not adaptive, the underlying rules are adaptive in that, in a natural environment in which bees are competing with many others, both rules generate an ideal-free distribution in which no bee could do better.

It would be interesting to consider what sorts of simple heuristics might produce impulsiveness and, more specifically, roughly hyperbolic delay-of-reinforcement gradients. Hyperbolic gradients are contrasted with the "rational" expectation of exponential decay (Read, 2004). Impulsiveness research studies the

balancing of rewards and delays, and we would like to draw a parallel with expected utility (EU) theory which concerns the balancing of rewards and probabilities, typically as revealed by whether subjects would be willing to take specified monetary gambles. The response to the well-documented violations of EU theory has been to design modifications such as Prospect Theory in which utility and probability functions are adjusted into more complex shapes. Prospect Theory is typically considered as an outcome model. An alternative is to design and test process models that would explain, rather than merely describe, the violations of EU. ABC (E. Brandstätter, G. Gigerenzer, R. Hertwig) has tested a heuristic much like Take The Best called the Priority Heuristic. This checks first the minimal gains (losses) and if their difference exceeds a preset aspiration level, search is stopped and the better gamble taken. If not, the probabilities of the minimal gains are checked, and then the maximal gains. Unlike EU and its modifications, the Priority Heuristic is lexicographic and does not require the animal to trade-off probabilities with gains or losses. This heuristic predicted 87% of the most popular choices in 260 monetary gambles from four published datasets; Cumulative Prospect Theory achieved only 77%. Might it be possible to find an equally simple heuristic than can generate the data that have been modelled with hyperbolic delay-of-reinforcement gradients?

3. Optimisation, optima and optimality modelling

Optimisation is also not a process model. In fact two phenomena are often confused under the term "optimise". One concerns the outcome of performing optimally. Like ecological rationality, this sense of optimise can only be defined with respect to both a currency and a specified environment. The other meaning of optimise is the process of finding an optimum. Often the term in this sense is applied to a method that tries to find the optimum but which might not get there (not all disciplines agree with this broadening of the term). Numerical analysis is full of different algorithms (cf. heuristics) for finding optima, including a variety of hill climbing algorithms as alluded to by Kemp. These and other optimisation algorithms may not succeed in finding global optima, often because they get stuck on a

local optimum. So observing non-optimal behaviour is no guarantee that the animal is not using optimisation. And observing optimal behaviour is no guarantee that the animal is not using simple heuristics.

Optimality models are concerned with predicting optima; the numerical methods of optimisation are incidental and of course those used by modellers are likely to be different to the way that an animal might arrive at an optimum. The utility of optimality models is a logically distinct issue to that whether animals behave optimally (Godfrey-Smith, 2001). Optimality models may be useful tools even if changing environments and cognitive constraints mean that animals are not unboundedly optimal. We sense three instances in the commentaries of a misunderstanding of optimality models.

(1) Optimality modelling is intended to understand the adaptive function of a behaviour in the environment in which the animal has evolved. The concept of ecological rationality emphasises that heuristics adapted to one environment may not be adaptive in another, so we should not generally expect optimality models to fit behaviour in strange laboratory environments in which animals have not evolved (unless we believe learning to be infinitely powerful). Indeed Stolarz-Fantino and Fantino provide a suite of excellent examples demonstrating how heuristics adapted to natural environments need not produce behaviours appropriate to laboratory environments. It is laboratory examples that Williams uses to support his claim of the failure of optimality models. He may well be right to use this evidence to criticise the expectations of other researchers that animals would behave optimally in the laboratory, but it is unnecessary to invoke cognitive limitations as an explanation.

(2) Mazur sees it as a criticism that optimality modelling might not predict features of learning a priori, but only account for them if they are known to the modeller in advance. In fact it is usually in the latter role that optimality modelling has proved useful in behavioural ecology. The aim is to test an adaptive explanation for a known behaviour, not usually to predict a previously unknown behaviour. For instance one of us has used optimality modelling to scrutinise various adaptive explanations of why there is a burst of bird song at dawn, a phenomenon of which humans must always have been aware (Hutchinson, 2002). Support for the explanation is strengthened if the optimality model makes further

independent predictions that fit the facts. An optimality model's most useful predictions are often of the range of circumstances under which one behaviour would be favoured rather than another. It would indeed be interesting to know when the detailed learning phenomena described by Mazur are adaptive and whether there is any correspondence to when they occur.

(3) Hirota and Sakagami are right that it may be difficult for human heuristics to select a performance measure that is biologically appropriate. However, they are mistaken in believing that biologists have not coped with similar problems. For instance, in biomechanics it may be uncertain whether to model a structure as having evolved to maximise stiffness or toughness for a given mass of material, and whether to resist loads in bending, compression or twisting (e.g. Currey, 2003). Finding a currency or combination of currencies which predicts the observed structure leads to an understanding of its function: this is the essence of optimality modelling. In behavioural ecology also there is often debate about which currencies most closely relate to fitness (e.g. Ydenberg et al., 1994; Hedenström and Alerstam, 1995). Just as with accuracy and frugality, two currencies may also be antagonistic in behavioural ecology (e.g. feeding rate and avoidance of predation: Houston et al., 1993). Two numerical techniques used by optimality modellers in such circumstances are Pontryagin's Maximum Principle and Stochastic Dynamic Programming.

4. Heuristic selection, learning and tuning

We have shown that if heuristics are to perform well, it is important that an appropriate one is used for each task. Selection of one heuristic from several that might be applied grades into processes concerned with learning heuristics and with learning parameters of heuristics (such as which cues to use and/or in which order). These issues were consistently commented on (Pitts, Sanabria and Killeen, Hirota and Sakagami, and Shettleworth). Part of the reason for this concern might be the attention given to "controlling variables" in behavioural analysis. Indeed this approach might profitably be applied to the selection of heuristics (Pitts), although, as with the heuristics themselves, we do not expect there to be a universal method of heuristic selection that can be studied in the laboratory

and that will be used in all circumstances outside it.

We very much agree that how heuristics are chosen is an important question that must be answered before a complete understanding is obtained, but we would maintain that partial ignorance of this aspect should not prevent us from addressing other issues such as ecological rationality. It is surely sensible (not a case of being "fooled"—Pitts) to draw common inferences about the adaptive function of heuristics that differ only in whether they have been learnt or hardwired. Similarly, biologists who, for instance, study the adaptive function of countershading (Ruxton et al., 2004), are not worried that the developmental background to this trait must vary widely between the many organisms (e.g. insects, fish, mammals, besides even military aircraft) in which it has independently evolved. Likewise, to answer one question of Sanabria and Killeen, it is of interest whether a heuristic is applied consciously or unconsciously, but this need not matter for the issue of ecological rationality on which ABC has so far concentrated.

A range of solutions to the problem of heuristic selection are conceivable. Cross and Jackson provide some nice examples of three alternatives: hardwiring in response to evolutionary experience, individual learning by trial and error, and something closer to calculation. We would add the possibility that humans may learn how to select a heuristic through being instructed by others. Hardwiring need not lack flexibility; the circuit could involve a sensitivity to a particular cue which would trigger different heuristics in different circumstances. Likewise if one has learnt painfully through trial and error that one heuristic is the best in a particular environment, we would expect that in different but similar circumstances certain similarities might trigger us to select the same heuristics from the start.

Hirota and Sakagami are right that simple heuristics can often be honed in performance as much by learning as by natural selection. Sanabria and Killeen wonder to what extent conventional learning processes are sufficient to explain our learning and selection of heuristics. Certainly these might play a role, although surely there are just too many potential cues and heuristics to have a good chance of finding a good combination unless the hunt is constrained. Shettleworth gives examples of why it is not best (and not done) to be equally sensitive to learn all cues. One problem with learning

from scratch is that learning processes effectively can become stuck on local optima and do not explore alternative cues and approaches. This was apparent in comparing heuristics for learning a good cue order for Take The Best from feedback of its success: if the best cues happened to start off low down the cue order they were never tried out and so remained there (Dieckmann and Todd, 2004).

Rieskamp and Otto (submitted for publication) provide a concrete example of how learning can guide the choice of heuristics. Subjects had to predict which of two firms would be more creditworthy given a number of quantitative cues such as "capital structure". A succession of such choices was presented to each subject over about 1 h, with immediate feedback given whether each decision was correct. What was specified correct depended on the environment. Subjects were each allocated to one of two environments, one non-compensatory in which Take The Best performed better than a weighted additive strategy, and one compensatory in which the reverse was true. The experimenters estimated which of these alternative heuristics was used from the decisions. When Take The Best was the more ecological rational, its usage increased through the trial (e.g. from under 30% to about 70%); in contrast, in the compensatory environment its usage decreased. Rieskamp and Otto fitted a reinforcement learning model to the data to describe how people learn to select the strategy that is more appropriate for an environment. Undoubtedly more such empirical examples of heuristic selection or learning would be highly desirable.

5. A confession and a route to salvation

One of us (JMCH) must now confess that he had only the haziest ideas about the field of behaviour analysis, despite having spent 15 years in post-doctoral research as a behavioural ecologist. More shocking still, he believes that he is typical. If behavioural ecologists become more interested in the mechanisms that enable

their adaptive strategies, and if behaviour analysts address how their results apply to natural environments, then we hope that this chasm in the study of animal behaviour can begin to be bridged.

References

Currey, J.D., 2003. How well are bones designed to resist fracture? J. Bone Miner. Res. 18, 591–598.

Dieckmann, A., Todd, P.M., 2004. Simple ways to construct search orders. In: Forbus, K., Gentner, D., Regier, T. (Eds.), Proceedings of the 26th Annual Conference of the Cognitive Science Society, Lawrence Erlbaum, Mahwah, NJ, pp. 309–314.

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

Gigerenzer, G., 2002. Reckoning with Risk: Learning to Live with Uncertainty. Allen Lane Penguin, London (USA edition: Calculated Risks: How to Know when Numbers Deceive You, Simon & Schuster, New York).

Godfrey-Smith, P., 2001. Three kinds of adaptationism. In: Orzack, S.H., Sober, E. (Eds.), Adaptationism and Optimality. Cambridge University Press, Cambridge, UK, pp. 335–357.

Hedenström, A., Alerstam, A., 1995. Optimal flight speed of birds. Philos. Trans. R. Soc. Lond. B 349, 471–487.

Houston, A.I., McNamara, J.M., Hutchinson, J.M.C., 1993. General results concerning the trade-off between gaining energy and avoiding predation. Philos. Trans. R. Soc. Lond. B 341, 375–397.

Hutchinson, J.M.C., 2002. Two explanations of the dawn chorus compared: how monotonically changing light levels favour a short break from singing. Anim. Behav. 64, 527–539.

Laland, K.N., 2001. Imitation, social learning, and preparedness as mechanisms of bounded rationality. In: Gigerenzer, G., Selten, R. (Eds.), Bounded Rationality: The Adaptive Toolbox. MIT Press, MA, pp. 233–247.

Read, D., 2004. Intertemporal choice. In: Koehler, D.J., Harvery, N. (Eds.), Blackwell Handbook of Judgement and Decision Making. Blackwell, Oxford, pp. 424–443.

Rieskamp, J., Otto, P.E., Adaptive strategy selection in decision making: the learning rule theory, submitted for publication.

Ruxton, G.D., Speed, M.P., Kelly, D.J., 2004. What, if anything, is the adaptive function of countershading? Anim. Behav. 68, 445–451.

Thuijsman, F., Peleg, B., Amitai, M., Shmida, A., 1995. Automata, matching and foraging behavior of bees. J. Theor. Biol. 175, 305–316.

Ydenberg, R.C., Welham, C.V.J., Schmid-Hempel, R., Schmid-Hempel, P., Beauchamp, G., 1994. Time and energy constraints and the relationships between currencies in foraging theory. Behav. Ecol. 5, 28–34.