

Testing the Ecological Rationality of Base Rate Neglect

Peter M. Todd*

*Center for Adaptive Behavior and Cognition
Max Planck Institute for Human Development
Lentzeallee 94, 14195 Berlin, Germany
ptodd@mpib-berlin.mpg.de

Adam S. Goodie**

**Department of Psychology
University of Georgia
Athens, GA 30602-3013 USA
goodie@egon.psy.uga.edu

Abstract

What simple learning rules can allow agents to cope with changing environments? We tested whether a rule that neglects base rates of events in the world — something that is usually considered irrational — could be as successful as Bayesian inference — the usual standard of rationality — in making cue-based predictions about events in time-varying environments. We focused on environments in which base rates change more frequently than cue accuracies, a condition that, we argue, is common in the real world. Five strategies (Bayesian, cue accuracy alone, adjusted cue accuracy, base rates alone, and a Least Mean Square learning rule) were compared across “lifetimes” of 10,000 predictions, in which base rates and cue accuracy independently changed every 10, 50, 100, 500, 1000, or 5000 events. The results confirmed that simple strategies that are typically deemed irrational (base rate neglect and its opposite, conservatism) can rival the typical standard of rationality, Bayesian combination of information, by producing ecologically rational decisions in appropriately varying environments.

1. Introduction

The best predictor of the future is the past. But what aspects of the past should be used to predict the events of the future? The information that could possibly be used includes the prevalence of particular past events — that is, their base rates — and cues that more or less reliably preceded past events. Given this information, making predictions based on combining both sources according to Bayesian updating is usually considered to be the best (most rational, even optimal) approach. However, in a wide array of prediction tasks, people have been found to underweight (or “neglect”) base rates and overweight situation-specific cues relative to Bayesian prescriptions. As an example of a situation in which people would be likely to neglect base rates (Nisbett & Ross, 1980), consider choosing between two car brands to buy, the first of which has a better long-term repair record. But then

you hear from a friend with that brand that their car just died on the highway — now which brand would you be likely to pick? Here, a single new piece of information should not overwhelm the evidence of cumulative repair records (base rates) — but in similar sorts of experimental tasks it (often) does. There are many manipulations that mitigate this effect (Koehler, 1996), but it is difficult to eliminate it entirely in many laboratory experiments. Indeed, in summing up a wide range of research, Bar-Hillel (1980) stated that “The genuineness, the robustness, and the generality of the base-rate fallacy [i.e., base rate neglect] are matters of established fact.”

Why do people (sometimes) neglect base rates in this way? More generally, why do people (as well as other species) systematically make decisions like these that are typically described as “irrational”? A large literature explores the conditions under which “reasoning errors” can be elicited (e.g., Kahneman, Slovic, & Tversky, 1982), while an opposing tradition demonstrates how such errors can be eliminated or substantially reduced (e.g., Gigerenzer, 1991). One view that strives to explain both sets of findings posits that human decision mechanisms are designed to work with specific structures of information in the environment. When these mechanisms are used in inappropriate environments, irrational decisions can result. But when they are applied to decision problems in appropriately-structured environments, humans can make decisions with *ecological rationality* — that is, appropriate choices that result from a match between a decision mechanism and the environment in which it is applied (Gigerenzer & Todd, 1999). In this paper, we argue that learning and decision mechanisms that produce a pattern of behavior usually held to be irrational, base rate neglect, can actually be ecologically rational by producing appropriate decisions in environments with appropriate structure. We demonstrate this match between decision strategies and environments through Monte Carlo simulations. Furthermore, we argue that the environment structure most suited for base-rate-neglecting mechanisms is actually common in the real world. Thus, the fact that people can be shown to inappropriately neglect base rates in laboratory experiments may be an indication that the experimental setups violate people’s reasonable expectations about environment structure, rather than

demonstrations of deep-seated irrationality.

As an example of a situation in which base rate neglect may be reasonable, Gigerenzer (1991) introduced a variation on the dying-car problem that casts some doubt on the universal applicability of the traditional Bayesian approach. Consider deciding whether to let your child play in the local river or in the nearby forest. Crocodiles have not been seen in the river for the past several years — establishing a safe base rate for choosing the river over somewhat risky tree climbing — but yesterday one attacked somebody at the river’s edge. Where should you let your child play today? In this case, it would seem the wiser strategy to emphasize the recent single case and discount (or, better, discard) the overall base rate.

The crocodile problem illustrates the importance of considering decisions in the context of the dynamic environment in which they are likely to occur. For the car-buyer, the base rates of brand reliability may be stable over the course of a few years, and a single breakdown might be merely an erroneous cue. For the parent, however, it is clear that the base rates of crocodile infestation can vary from year to year, while a recent report of a crocodile attack remains a reliable cue to the present situation. The cue of yesterday’s crocodile attack may serve not only its traditional role in the likelihood ratio of Bayes’ formula, but may also indicate that the prior probabilities — the base rates — have changed.

In light of such considerations, it has been proposed (Gigerenzer, 1991; Goodie & Fantino, 1999) that in environments where base rates change more than cue accuracies¹, neglecting base rates when making predictions may not impair performance to any great extent. Thus, observations of “irrational” base rate neglect may indicate an otherwise-ecologically rational mechanism is being used in an environment for which it is not intended. In this paper we test this idea empirically, comparing in simulation the performance of 5 prediction strategies derived from the literature, including partial and complete base rate neglect, Bayesian integration, base rate use alone and connectionist learning, in a range of environments with varying cue accuracies and event base rates.

2. Past work on ecologically rational learning

The question of which particular learning strategies are ecologically rational in which environments is closely related to the question of when (i.e., in what environmental circumstances) different forms of learning will evolve. While there is wide agreement that rapidly changing environments will be favorable to the evolution of learning in general (e.g., Belew, 1990; Miller & Todd, 1990; Stephens, 1991), there is relatively little exploration to date of specific environment structures that favor the evolution of particular types of learning. Todd & Miller (1991) evolved simple recurrent neural networks that exhibited the ability to sensitize and habituate in “clumpy” environments where important stimuli were encountered in

¹We use the term cue accuracy to mean Bayesian likelihoods, in other words $p(\text{cue} \rightarrow \text{outcome})$.

(more or less noisy) patches over time. Thus in effect the networks experienced changing base rates, but there were no explicit cues (other than the presence or absence of the stimuli) to use along with base rates, greatly limiting the range of possible learning strategies that could be employed. Here we extend this early work by explicitly comparing a set of clearly defined (rather than messily evolved) learning strategies deployed in a wider range of environment structures to determine when they fit, and thus produce ecological rationality.

McKenzie (1994) took a similar direct approach by simulating intuitive strategies and more traditional rational models making decisions in a broad range of environments. He found high correlations in performance between the intuitive strategies and the rational models in many instances, indicating the ecological rationality of the former. (Gigerenzer & Hoffrage, 1995, extended the catalog of intuitive near- and non-Bayesian strategies that people use when faced with Bayesian decision tasks.) Our work in this paper expands on McKenzie’s efforts by exploring the important case of strategies tuned to changing, rather than fixed, environments.

3. Simulating cue and outcome learning

To test the efficacy of different prediction strategies in different settings, we created a simulation that can be viewed as comparing the behavior of organisms following various prediction strategies in particular environments. The “lifetime” of each simulated organism consisted of 10,000 occasions on which either of two events, E_0 or E_1 , could occur, preceded in each case by one of two cues (or two states of one cue, e.g. absent or present), C_0 or C_1 . The organism used its particular strategy to predict the event after “perceiving” the preceding cue, and the success of each strategy was measured as the average number of events correctly predicted over 50 lifetimes.

3.1 The environments

To create environments in which event base rates and cue accuracies can vary over time, we specified two parameters, the base rate change interval (*BRCI*) and the cue accuracy change interval (*CACI*). Each parameter was assigned intervals of 10, 50, 100, 500, 1000, or 5000 events, yielding 36 environments. For each organism’s lifetime, an initial base rate and cue accuracy were chosen randomly from a uniform distribution between 0.0 and 1.0. Then, after *BRCI* events, the base rate was changed to a new value chosen at random from the same interval, and after *CACI* events, a new cue accuracy was similarly selected. So for example an environment characterized by a *BRCI* of 50 and a *CACI* of 1000 has base rates that change rapidly while cue accuracies remain relatively stable — the kind of situation in which base rate neglect might be adaptive.

The current base rate and cue accuracy values controlled the generation of the events and cues an organism would experience on a given trial. First, event E_0 or E_1 was selected

randomly with a probability distribution defined by the current base rate. For example, if the base rate were .7, E_0 would be chosen with 30% probability and E_1 with 70% probability. Second, based on the selected event, a cue would be generated according to the cue accuracy. (Note that in these simulations, we make the simplifying assumption that $p(C_0|E_0) = p(C_1|E_1)$.) For example, if E_1 was chosen and the cue accuracy was .4, then C_1 would be generated with 40% probability, and otherwise C_0 would be generated. In this way, 50 lifetime sequences of 10,000 cues and events were generated for each of the 36 *BRCI/CACI*-specified environments. Each strategy described in the next section was tested on the same set of 50 sequences per environment.

3.2 The strategies

We compared the performance of five strategies in two classes: four that kept track of cue-event pairs in a decaying memory table, and one that learned connection strengths between cues and events. Each strategy computed the probability p of event E_1 given the observed cue (C_0 or C_1) and its knowledge of the past, and then predicted E_1 if $p > 0.5$ and E_0 if $p \leq 0.5$.

3.2.1 Contingency table-based strategies

Four of our strategies are based on memory for the frequencies of cue-event pairings as represented in the 2×2 table shown in Figure 1. The most straightforward application of this approach would imply a full and perfect memory of M events, with complete amnesia for all observations $M + 1$ events ago and earlier. To use a more realistic model of memory that incorporates the assumption that old information is less relevant than recent information, we employed a table in which stored events decay exponentially as further events occur. For instance, entry a in the table is updated as follows: $a_{t+1} = f \cdot a_t + 1$ if the cue-event pair C_0, E_0 was just seen, and $a_{t+1} = f \cdot a_t$ otherwise, where f is the decay rate parameter ($0 < f < 1$). In this way, $a, b, c,$ and d maintain a weighted count of the past event-cue pairs seen, with the most recent observations getting the most weight. To paraphrase Douglas MacArthur², old events never depart from memory; they just fade away. Each of the four table-based strategies we investigated then computed its prediction on the basis of the values $a, b, c,$ and d , as described below and summarized in Table 1. We varied the memory decay rate f , but to allow memory length to be compared more directly to rates of environment change, we give our results in terms of “half-life,” which is the number of trials it takes a given event-cue pair entered into memory initially with a strength of 1.0 to decay to a strength of 0.5 (half-life $\approx -1/\log_2 f$).

1. Bayesian. This strategy consists of Bayesian integration of the base rates and cue accuracy, defined as the conditional probabilities of E_1 given the observed cue. This strat-

| | | Cue | | |
|-------|-------|-------|-------|-------------|
| | | C_0 | C_1 | |
| Event | E_0 | a | b | a+b |
| | E_1 | c | d | c+d |
| | | a+c | b+d | M = a+b+c+d |

Figure 1: The structure of the memory for strategies 1 to 4. A decaying memory trace was kept for the four cells $a, b, c,$ and d , and different strategies used some or all of these values.

egy employs all four pieces of available information, and is usually considered the benchmark for rationality in this domain (Castellan, 1977).

2. Cue accuracy alone (base rate neglect). This strategy considers only the probability that the cue matches the upcoming event. By neglecting base rate information, this strategy should be outperformed by the Bayesian strategy 1 in environments where base rates are stable and useful, but may gain an advantage when base rates change quickly.

3. Adjusted cue accuracy. Absolute base rate neglect as embodied in strategy 2 has seldom been reported in the literature (Koehler, 1996; but see Gigerenzer & Hoffrage, 1995, for instances). Rather, humans and other animals seem to incorporate base rate information partially in their decisions, as is modeled by this strategy. The cue-based probability used in strategy 2 was adjusted by a small amount (5 percentage points) in the direction prescribed by the current base rate. This strategy is consistent with the substantial but incomplete base rate neglect observed by Goodie and Fantino (1995; 1999).

4. Base rates alone. The complement to using cue accuracy exclusively is using base rates exclusively. Where the former neglects base rates entirely, this strategy neglects cue-based information entirely, employing only the probability with which each event occurs. This strategy is similar to an extreme form of conservatism in decision making (Edwards, 1982). While conservatism has been described as erroneous reasoning, this strategy should perform well when cue accuracy changes frequently.

3.2.2 The connectionist learning strategy

The fifth strategy arises from a connectionist model of categorization that has been applied to base rate phenomena. This strategy dispenses with directly counted memories of events and instead learns continuously shifting association strengths between experienced cues and events. A set of input nodes receives activation according to the current cue, and an output node makes an event prediction based on the summed input activation weighted by the learned connection weights between input and output nodes. Learning takes place after

²Address to a Joint Meeting of Congress, April 19, 1951.

each cue-event pair is seen, adjusting the connection weights to minimize the difference between predictions and actual events.

5. Least Mean Squares (LMS) learning rule. Gluck and Bower (1988) trained people to diagnose fictitious diseases from symptoms in a situation analogous to the cue-to-event prediction we consider here. They found base rate neglect that could be modeled with a simple neural network trained according to the LMS learning rule. Our strategy used a similarly trained network with three input nodes: one activated by C_0 , one by C_1 , and a third that was always activated, simulating background cues. We followed Gluck and Bower’s formulas for updating activation weights and converting them to response probabilities.

| | prediction after C_0 | prediction after C_1 |
|--------------------------|--|--|
| 1. Bayesian | $R(1 - \frac{a}{a+c})$ | $R(1 - \frac{d}{b+d})$ |
| 2. Cue accuracy | $R(1 - \frac{a+d}{M})$ | $R(\frac{a+d}{M})$ |
| 3. Adjusted cue accuracy | $R(1 - \frac{a+d}{M} + .05)$ if $\frac{a+b}{M} < .5$ $R(1 - \frac{a+d}{M} - .05)$ if $\frac{a+b}{M} > .5$ | $R(\frac{a+d}{M} + .05)$ if $\frac{a+b}{M} > .5$ $R(\frac{a+d}{M} - .05)$ if $\frac{a+b}{M} < .5$ |
| 4. Base rates | $R(\frac{c+d}{M})$ | $R(\frac{c+d}{M})$ |

Table 1: Formulas for calculating the predictions of the contingency-table-based strategies (1-4) from the information stored in the 2×2 table shown in Figure 1. $R(x)$ denotes the function of rounding a number to the nearest integer. Because all values represented here are bounded by 0 and 1, all rounded values are 0 or 1, corresponding to predicting E_0 or E_1 , respectively. Some values to be rounded are subtracted from 1 because, for example, the cue accuracy strategy should predict E_0 when $a + d$ is close to 1, not when it is close to 0.

3.3 Predictions and Questions

These 5 strategies were compared by simulating their predictions over 50 lifetimes of 10,000 cue-event pairs in each of 36 different environments. The most immediate question addressed by this comparison is whether base rate neglect (strategy 2) can successfully compete with Bayesian integration (strategy 1) when base rates change more frequently than cue accuracy. Of course, if the Bayesian strategy has accurate knowledge of the base rates and cue accuracies in the current environment at every moment, it will have an advantage in using this knowledge. But such knowledge does not come from a guidebook — it must be acquired through experience. If base rates change so frequently that they are often estimated incorrectly, then a strategy that uses these estimates may have

little advantage over a strategy that ignores base rate information. We expected this to be the case when the Bayesian strategy attempted to incorporate rapidly changing base rates. Particularly when cue accuracy changed seldom, we expected base rate neglect strategies (2 and 3) to perform about as well as Bayesian integration (strategy 1).

We correspondingly expected that “cue-accuracy neglect” (strategy 4) would rival Bayesian integration (strategy 1) in environments with rapidly-shifting cue accuracies and stable base rates. Finally, we were not sure how a connectionist strategy would fare in this comparison; in principle, it could learn and use both base rate and cue accuracy information, like Bayesian integration. But because such models have been used to predict (correctly) base rate neglect in comparable situations, we thought strategy 5 might behave similarly to strategies 2 and 3.

4. Comparing rule performance in different environments

The prediction performance of all 5 strategies³, using a memory half-life of 100 cue-event pairs for strategies 1-4, is shown in Figure 2. Averaged across all 36 environments, the Bayesian integration strategy proved the most successful, making correct predictions for about 73% of the cue-event pairs (see Figure 2, left panel). The adjusted cue accuracy strategy (3), which must keep track of three of the cells in the 2×2 table, combined this information in a slightly less effective manner, reaching about 69% correct predictions. This was followed closely by the LMS strategy with 68%. The two simplest strategies, using only two of the cells in the 2×2 table (cue accuracies or base rates alone), each scored about 66% correct overall.

However, when environments with different structures are examined separately, the ordering of successful strategies changes systematically and dramatically. First, consider those environments where base rates are much less stable than cue accuracies, defined as those where the interval between cue accuracy changes is at least ten times as long as the interval between base rate changes ($CACT \geq 10 \cdot BRCI$). Ten of the 36 environments have this structure. Averaged across these environments (see Figure 2, middle panel), Bayesian integration still makes correct predictions most often, even exploiting this environment structure to improve performance slightly to 74%. However, base rate neglect now nearly matches Bayesian performance, with 72% correct. Including the base rate adjustment (strategy 3) improves performance by another percentage point (73% correct). The LMS rule’s performance is reduced by the rapidly changing base rates to 65%. Finally, strategy 4, in relying solely on base rates,

³We performed a similar comparison across these three environment structures for two additional types of strategies: a set corresponding to strategies 1-4 that used an exact count of cue-event pairs in a sliding memory window rather than a decaying memory; and a set corresponding to strategies 1-5 that predicted outcome events using probability matching rather than exclusively choosing the more likely event. In both cases, the patterns of results closely resembled those reported here.

Overall Performance

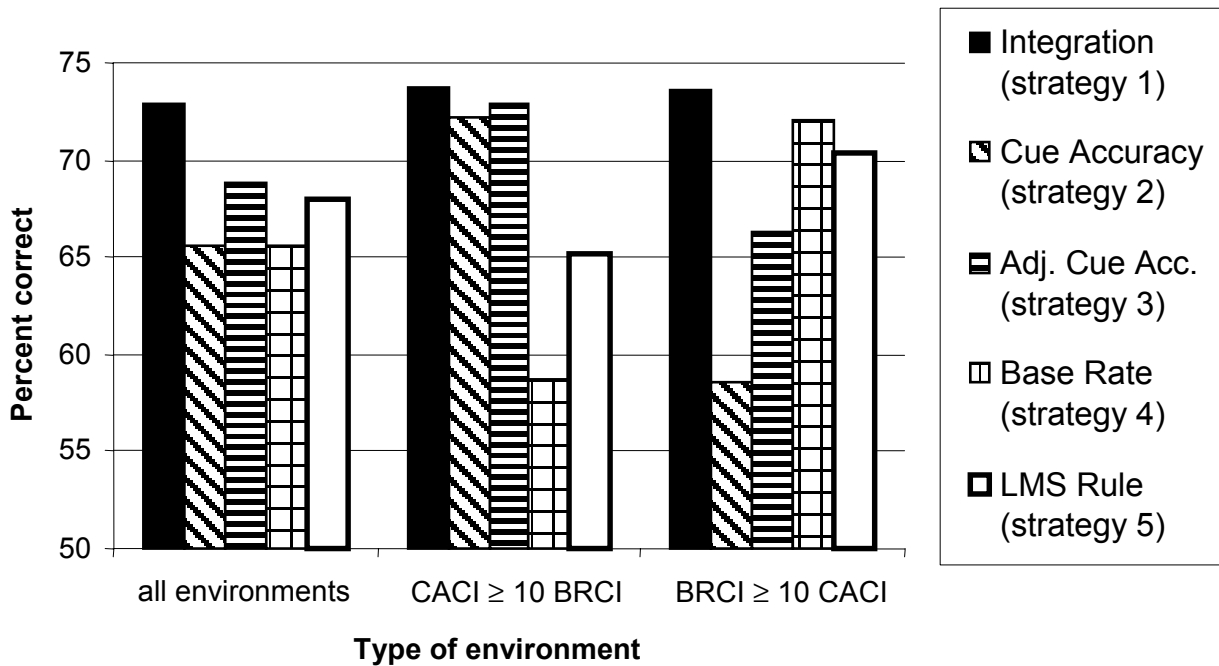


Figure 2: The performance of the five prediction strategies averaged across all 36 environments (left panel), across those environments with stable cue accuracies and rapidly varying base rates (middle panel), and across environments with stable base rates and rapidly varying cue accuracies (right panel).

not surprisingly does worst of all, but is still able to extract enough information to make correct predictions 59% of the time.

In environments where cue accuracies are much less stable than base rates ($BRCI \geq 10 \cdot CACI$; Figure 2, right panel), we find a mirror image of these results. Base rates alone (strategy 4) now nearly matches Bayesian performance (72% vs. 74%), while base rate neglect performs worst (59%). Partial base rate neglect (strategy 3) achieves an intermediate accuracy (66%). The LMS rule also benefits from stable base rates (improving to 70%).

Because our primary question was when base rate neglect would approach the accuracy of the Bayesian strategy, the difference in performance between these two strategies is portrayed across all environments in Figure 3. When cue accuracies changed slowly and base rates changed rapidly (front left corner), base rate neglect was about as accurate as Bayesian integration. As cue accuracies also changed more and more rapidly (moving along the front edge toward the right), the two strategies continued to perform similarly, in part because all strategies fared less well in environments where both cue accuracies and base rates changed rapidly (front right corner). Only when base rates became much more stable ($BRCI \geq 500$) did a large difference between the two strategies appear, particularly when cue accuracies were less reliable than base rates (rear right corner).

In comparing the LMS learning rule with the contingency-table strategies, three unpredicted but robust effects emerged.

First, the LMS rule did a better job of capitalizing on stable base rates than on stable cue accuracies (see Figure 2, difference between right and middle panels, respectively). Second, this strategy performed better than any other strategy in the most variable environments ($BRCI, CACI = 50$; see Figure 4). Third, however, it lags behind other strategies in environments of frequent base rate change where base rate neglect does well (Figure 2, middle panel), perhaps because it does relatively poorly at taking advantage of stable cue accuracies. This is at first surprising in view of Gluck and Bower's (1988; see also Shanks, 1990) demonstration that their model using the LMS rule produces base rate neglect. This paradox may be clarified by noting the distinction between base rate neglect as an empirical phenomenon, as the term is most commonly used, and base rate neglect as a decision process, as was employed here — the LMS rule can produce base rate neglect phenomena in some circumstances, but our results make it clear that it does so in a way that differs from the cue-accuracy strategy (2).

4.1 Effects of memory length

Do these effects hold with other assumptions about the length of memory for cue-event pairs? We tested the robustness of our results to changes in memory length by comparing the performance of strategies 1-4 with half-lives of 20, 50, 100, 200, 500, 1000, 2000, and 5000 events. We used four environments combining slow and fast

Difference in Performance Between Bayesian integration (strategy 1) and base rate neglect (strategy 2)




Figure 3: Comparison of the performance of Bayesian integration and base rate neglect across all 36 environments specified by *BRCI* and *CACI*. The z-axis shows the difference in number of correct predictions (out of 10,000 predictions made) between the two strategies (Bayes minus base rate neglect).

changes: $BRCI/CACI = 1000/1000, 1000/50, 50/1000,$ and $50/50$. The results appear in Figure 4.

Memory length has a strong effect on overall performance: Longer memory traces almost always impair performance. This is because longer memories blur current and outdated base rate and cue accuracy information. Furthermore, the main patterns seen before remain robust across different memory lengths: When base rates change much more rapidly than cue accuracy ($BRCI = 50, CACI = 1000$), base rate neglect (strategy 2) performs nearly as well as Bayesian integration (strategy 1) for all but the shortest memories. Likewise, when cue accuracy changes more rapidly than base rates ($BRCI = 1000, CACI = 50$), base rate use (strategy 4) performs nearly as well as Bayesian integration for all but the shortest memories. Because short memory lengths do not seem likely (in light of, e.g., Goodie & Fantino’s 1999 learning studies), the ability of simple strategies to match Bayesian performance in appropriate environments appears robust under reasonable assumptions of memory length.

5. Forms of environmental variation

Overall, these results show that simple predictive strategies matched to environments with a specific structure can rival “optimal” models in prediction performance. We compared the performance of a variety of decision strategies across a range of environments, and found that although some of these strategies violate traditional norms of rationality, for instance by ignoring or underweighting base rates, they produce good decisions in environments to which they are matched, and thereby yield ecological rationality (Gigerenzer & Todd, 1999).

But *do* base rates change more often than cue accuracies

in general? We believe they do, because base rates of natural events are usually caused by complex and interacting factors, while good cues often bear a relatively simple, and therefore stable causal relation to the events they predict. As a simple example, consider the problem of predicting picnic weather. The base rate of rain on a particular day in a particular place is complexly determined by a number of time-varying factors that are often difficult to evaluate, such as approaching fronts, humidity, microclimate, and time of year, making accurate assessment of the prior probability difficult and time-consuming. However, these multifarious causal variables often combine to create not only rain, but also rain clouds. Dark clouds are a good cue to impending rain, and because they are mostly caused by the same forces, the likelihood of storm clouds when rain is imminent is expected to stay constant, even as the base rates of both rain and rain clouds may vary widely. Similarly, incidence of measles may vary from time to time, or from one locale to another, but red spots are an accurate cue in each situation. Birth rates are different in different countries and in different years, but pregnancy tests remain equally reliable. Another way of thinking of this is that when things in the world are clumped in space or time, organisms will experience changing base rates of those things as they move through space or time. Just how clumped various things are compared to how stable (over space and time) their cues are will determine when base rates and cue accuracies are informative.

In contrast, we can ask what environments exist where cue accuracies vary more rapidly than base rates. In such environments, we could expect people to rely relatively more on base rate information, and thus not show (strong) base rate neglect. Situations in which cues are socially constructed, and so may vary rapidly, are one candidate for rapidly varying cue accuracies, especially when there are conflicts of interest involved between signalling and receiving individuals. Similarly, when entering a new arbitrarily constructed situation (as social environments can be), cue accuracies may be expected to change more than base rates.

Goodie and Fantino (1995, 1996) found evidence that people may expect such changes in cue accuracy, but not base rates, when encountering novel situations. In two sets of experiments they created conditions of direct experience in which, unlike in past studies, participants were not required to process or comprehend any statistical information, and did not have to report any quantitative judgments. First, they asked participants to predict events consisting of the occurrence of blue or green rectangles on the basis of cues that were either other blue or green rectangles (Goodie & Fantino, 1995) or the words “blue” or “green” (Goodie & Fantino, 1996, experiment 2). In both cases, participants showed standard base rate neglect, behaving as though they expected the cues to be stable, accurate predictors while the base rates could be variable and hence not useful. This is understandable in light of the facts that, on the one hand, we commonly match colors and match names to colors (thus experiencing

reliable cue accuracies), and on the other hand we see some colors more often than others in different situations (thus experiencing variable base rates). Second, when the cues were instead unrelated to the events, such as vertical or horizontal lines preceding the blue or green rectangles, base rate neglect disappeared entirely (Goodie & Fantino, 1996, Experiment 1), indicating that participants were now using a strategy consistent with the belief that both event base rates and cues could be important in this novel prediction setting.⁴

6. Implications and conclusions

While the events and cue-event associations of one's past are an important guide to the future, knowledge of the past can prove untrustworthy if it springs from processes that are no longer in place. And in the real world, change is constant. The changing seasons, the shifting populations of interacting predator and prey (or parasite and host) species, the new surroundings encountered by migrating creatures, all contribute to reducing the sameness of past and future (Potts, 1996).

With the present simulations we have demonstrated that a supposedly normative, Bayesian strategy is often little more successful than less "rational" strategies at predicting the events in particular changing environments. Thus, the fact that people can be made to act "irrationally" in underusing base rates (or conversely being too conservative in changing them) appears to say more about the environment structure that people expect (whether via learning or evolution) to encounter than about the purported shortcomings of our cognitive abilities. While there are several extensions to this work that must be explored before definitive statements can be made (e.g., eliminating the constraint of equal accuracies for both cues, using more extreme base rates and cue accuracies — possibly bimodally distributed — in the environments, and testing the efficacy of a wider range of strategies), our results here indicate that observed human behavior in many learning situations can be ecologically rational, rather than irrational.

In addition to the near-Bayesian performance of strategies that neglect base rates under presumably common shifting environmental conditions, two more of our findings bear on questions of standards of rationality. First is the finding that longer memories lead to poorer performance. It is perhaps surprising when more information (here stored in memory) results in reduced performance, but such "less is more" effects may not be so uncommon (e.g., Goldstein & Gigerenzer, 1999). This instance presumably arises from a compromised ability to respond rapidly to changes in the environment.

The other notable outcome is the performance of the LMS rule relative to the Bayesian strategy: sometimes better,

⁴Furthermore, when participants in the familiar matching situation were allowed to acquire a great amount of experience about the particular experimental environment — 1600 predictions — they began by neglecting event base rates as indicated above. Over time, though, their performance changed in a way that indicated the participants had learned that within this context a new and stable set of base rates held and hence a strategy incorporating both base rates and cues could again be used.

sometimes worse. Both LMS learning (Widrow & Hoff, 1960) and Bayesian decision making (Tversky & Kahneman, 1982) are widely proclaimed as optimal — but only under certain conditions. Indeed outside those conditions, as in the environments explored here, the two strategies do different things. How then can we choose an appropriate norm of rationality to use as a benchmark when judging the performance of difference decision mechanisms (and of human decision making; see Hertwig & Todd, 2000)? One could argue that whichever approach performs better is the truly rational model. But the fact that Bayesian integration does better in some environments while the LMS rule does better in others shows that norms, like the decision mechanisms themselves, must be chosen relative to a particular environment. Thus, the structure of the environment proves to be the final arbiter of ecological rationality.

Acknowledgements

We thank Valerie M. Chase, Gerd Gigerenzer, and Ralph Hertwig for detailed comments on earlier versions of the text.

References

- Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, 44, 211-233.
- Belew, R.K. (1990). Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems*, 4, 11-49.
- Castellan, N. J. (1977). Decision making with multiple probabilistic cues. In N. J. Castellan, D. P. Pisoni, & G. R. Potts (Eds.), *Cognitive theory*, Vol. 2 (pp. 117-147). Hillsdale, NJ: Erlbaum.
- Edwards, W. (1982). Conservatism in human information processing. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- Gigerenzer, G. (1991). How to make cognitive illusions disappear: Beyond "heuristics and biases". *European Review of Social Psychology*, 2, 83-115.
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instructions: Frequency formats. *Psychological Review*, 102, 684-704.
- Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In G. Gigerenzer, P. M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 3-34). New York: Oxford University Press.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227-247.
- Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 37-58). New York: Oxford University Press.

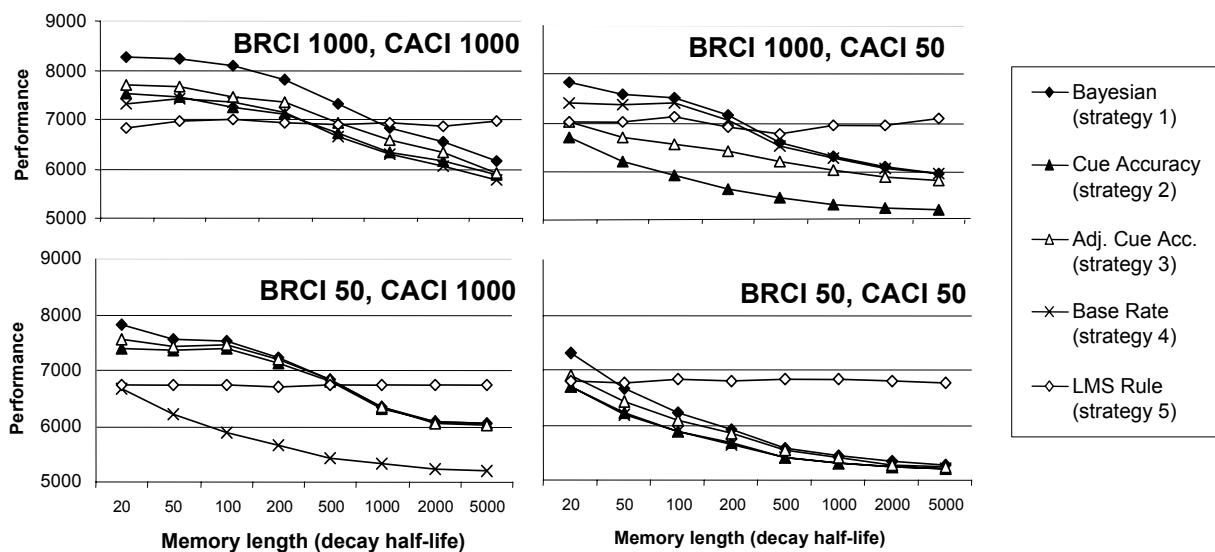


Figure 4: Performance (number of correct predictions out of 10,000) of strategies 1 to 4 with different memory decay parameters, tested in four environments characterized by different rates of base rate and cue accuracy change. The performance of the LMS rule (which is not affected by variations in the decay parameter) is shown in each case for comparison.

Goodie, A. S., & Fantino, E. (1995). An experientially derived base rate error in humans. *Psychological Science*, 6, 101-106.

Goodie, A. S., & Fantino, E. (1996). Learning to commit or avoid the base rate error. *Nature*, 380, 247-249.

Goodie, A. S., & Fantino, E. (1999). What does and does not alleviate base rate neglect under direct experience. *Journal of Behavioral Decision Making*, 12, 307-335.

Hertwig, R., and Todd, P. M. (2000). Biases to the left, fallacies to the right: Stuck in the middle with null hypothesis significance testing. *Commentary on Krueger on social-bias. Psycholoquy* 11(28). <http://www.cogsci.soton.ac.uk/psycbin/newpsy?11.028>

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.

Koehler, J. J. (1996). The base rate fallacy reconsidered: Descriptive, normative and methodological challenges. *Behavioral and Brain Sciences*, 19, 1-53.

Lovett, M. C., & Schunn, C. D. (1999). Task representations, strategy variability and base rate neglect. *Journal of Experimental Psychology: General*, 128, 107-130.

McKenzie, C. R. M. (1994). The accuracy of intuitive judgment strategies: Covariation assessment and Bayesian inference. *Cognitive Psychology*, 26, 209-239.

Miller, G.F., & Todd, P.M. (1990). Exploring adaptive agency I: Theory and methods for simulating the evolution of learning. In D.S. Touretzky, J.L. Elman, T.J. Sejnowski, and G.E. Hinton (Eds.), *Proceedings of the 1990 Connection-*

ist Models Summer School (pp. 65-80). San Mateo, CA: Morgan Kaufmann.

Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall.

Potts, R. (1996). *Humanity's descent: The consequences of ecological instability*. New York: William Morrow.

Shanks, D. R. (1990). Connectionism and the learning of probabilistic concepts. *Quarterly Journal of Experimental Psychology*, 42, 209-237.

Stephens, D.W. (1991). Change, regularity, and value in the evolution of animal learning. *Behavioral Ecology*, 2, 77-89.

Todd, P.M., & Miller, G.F. (1991). Exploring adaptive agency III: Simulating the evolution of habituation and sensitization. In H.-P. Schwefel and R. Maenner (Eds.), *Proceedings of the First International Conference on Parallel Problem Solving from Nature* (pp. 307-313). Berlin: Springer-Verlag.

Tversky, A., & Kahneman, D. (1982). Evidential impact of base rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press

Widrow, B., & Hoff, M. E. (1960). Adaptive switching circuits. *IRE WESCON Convention Record*. pp. 96-104. Reprinted in J. A. Anderson and E. Rosenfeld (Eds.), *Neurocomputing: Foundations of research* (pp. 126-134). Cambridge, MA: MIT Press.