

processes and beta prior distributions for the binomial probability parameters, natural sampling results in a mathematically nice property: The posterior distribution turns out to be also a beta distribution, and, most important in the present context, the distribution does not depend on the between-group frequencies (base rates) but only on the frequencies within each group (e.g., hit and false alarm rates). As the between-group frequencies are estimators of base-rate probabilities, this result describes a situation in which the base rates are irrelevant. Under properly defined natural sampling conditions, base-rate neglect is rational. The result was generalized to multinomial sampling in combination with Dirichlet priors (Kleiter & Kardinal 1995) and used to propagate probabilities in Bayesian networks (Kleiter 1996). Recently, Hooper (2007) has shown that some claims about the generality of beta posteriors in Bayesian networks made in my 1996 paper are only asymptotically valid.

If base rates are irrelevant in a “normative” model, then base-rate neglect in psychological experiments is not necessarily an error but may be rational. If Bayes’ Theorem is written in “frequency format,” even elementary school math shows that the base rates in the numerator and in the denominator get cancelled when the within-group frequencies add up to the between-group frequencies. This property fitted extremely well within Gigerenzer’s approach. In the early 1990s when Gerd Gigerenzer was at Salzburg University, during one of the weekly breakfast discussions held among Gigerenzer and members of his group, the mathematical result of base-rate cancellation was communicated and it was immediately taken up and integrated into his work. Natural sampling requires random sampling, additive frequencies in hierarchical tree-like sample/subsample structure (i.e., complete data), and a few more properties that belong to the statistical model. The notion of “natural frequencies” seems, in addition, to involve sequential sampling and thus acquires an evolutionary adaptive connotation.

The additivity in natural sampling goes hand in hand with the subset structure, the favorite explanation in the target article. The close relationship between natural sampling and the subset structure may have led to a confounding of the two in the past. If frequencies (and not subset structures) are the cause of facilitation effects, then critical experiments should investigate non-natural sampling conditions (Kleiter et al. 1997). Frequencies should still have a facilitating effect. Unfortunately, instead of non-natural sampling conditions, often “single-case probabilities” are taken for comparison to demonstrate the base-rate facilitation with natural sampling conditions.

How common are natural sampling conditions in everyday life? I have severe doubts about the ecological validity and the corresponding evolutionary adaptive value. From the perspective of ecological validity, it is important that the base-rate neglect has often been demonstrated for categories with low prevalence, such as rare diseases. Consequently, the prevalence of base-rate neglect will also be low. Base-rate effect certainly depends upon the actual numbers used in the experiments, a property not discussed in B&S’s review.

The cognitive system of an intelligent agent capable of uncertainty processing and judgment requires competence in at least six domains. (1) Perception and processing of environmental information, such as numerosity, cardinalities of sets, relative frequencies, descriptive statistics of central tendency, variability, and covariation. (2) Understanding of randomness, of not directly observable states, of alternatives to reality and hidden variables, of the non-uniformities in the environment, and of the limited predictability of events and states. (3) Introspection of one’s own knowledge states, and weighting and assessing one’s own incomplete knowledge by degrees of beliefs (subjective probabilities). (4) An inference engine that derives conclusions about the uncertainty of a target event from a set of uncertain premises. Typical inference forms are diagnosis, prediction, or explanation. The conclusions often concern single events. The probabilities can be precise or imprecise (lower and upper probabilities,

or second order probability distributions). Recently, classical deductive argument forms have also been modeled probabilistically (Oaksford & Chater 2007; Pfeifer & Kleiter 2005). (5) Modeling functional dependencies/independencies which are basic to causal reasoning. (6) Understanding of the knowledge states of other persons – a prerequisite for the effective communication of uncertainty in social settings.

Many base-rate studies present frequency information (belonging to item [1] in the list given above) and observe whether the subjects use “Bayes’ Theorem” as an inference rule (belonging to item [5]). Bayes’ Theorem degenerates to a rule for cardinalities, formulated not in terms of probabilities but in terms of frequencies (see Note 2 in the target article). This can of course be done, but we should be aware that we are dealing with the most elementary forms of uncertain reasoning, not involving any of the other items listed above. Moreover, if the response mode requires frequency estimates and not the probabilities of single events, another important aspect of uncertain reasoning is lost. If subjects are poor in the judgment of single event probabilities they have an essential deficit in uncertainty processing.

Conditional events and conditional probabilities are at the very heart of probability theory. Correspondingly, the understanding of conditional events and conditional probabilities should be central to investigations on human uncertain reasoning. Considering base-rate tasks in natural sampling conditions alone, misses this point completely. The B&S structural subset hypothesis shows that conditional probabilities are not needed in this case, and that structural task properties are the main cause of facilitation effects.

Dual concerns with the dualist approach

DOI: 10.1017/S0140525X0700180X

David A. Lagnado and David R. Shanks

Department of Psychology, University College London, London WC1E 6BT, United Kingdom.

d.lagnado@ucl.ac.uk d.shanks@ucl.ac.uk

Abstract: Barbey & Sloman (B&S) attribute all instances of normative base-rate usage to a rule-based system, and all instances of neglect to an associative system. As it stands, this argument is too simplistic, and indeed fails to explain either good or bad performance on the classic Medical Diagnosis problem.

Barbey & Sloman (B&S) claim that an associative system is responsible for errors in a range of probabilistic judgments. Although this is plausible in the case of the conjunction fallacy (where a similarity-based judgment substitutes for a probability judgment), it is less applicable to the Medical Diagnosis base-rate problem. What are the automatic associative processes that are supposed to drive incorrect responses in this case? Respondents reach incorrect solutions in various different ways (Brase et al. 2006; Eddy 1982), many of which involve explicit computations. Indeed, the modal answer is often equal to one minus the false positive rate (e.g., 95% when the false positive rate is 5%). This clearly involves an explicit calculation, not the output of an implicit process. Thus, errors can arise from incorrect application of rules (or application of incorrect rules), rather than just brute association.

The key point here is that base-rate neglect in the Medical Diagnosis problem provides little evidence for the exclusive operation of an implicit associative system. Indeed, it is arguable that adherents of classical statistics are guilty of similar base-rate neglect in their reliance on likelihood ratios (Howson & Urbach 2006). Presumably this is not due to an implicit associative system, but is based on explicit rules and assumptions.

What about the claim that the rule-based system is responsible for correct responses in frequency-based versions of the task? This hinges on the idea that representing the task in a frequency format alerts people to the relevant nested set relations, and thus permits the operation of the rule-based system. In one sense, this is trivially true – those participants who reach the correct answer can always be *described* as applying appropriate rules. But what are these rules? And what is the evidence for their use, as opposed to associative processes that could also yield a correct response?

B&S explicitly block one obvious answer – that once people are presented with information in a format that reveals the nested set structure, they use a simplified version of Bayes' rule to compute the final solution. The authors' reasons for rejecting this answer, however, are unconvincing. The cited regression analyses (Evans et al. 2002; Griffin & Buehler 1999) were performed on a different task. And they were computed on grouped data, so it is possible that those who answered correctly did weight information equally. Furthermore, it is wrong to assume that a Bayesian position requires equal weighting – in fact, a full Bayesian treatment would allow differential weights according to the judged reliability of the sources.

More pertinently, if people are not using the frequency version of Bayes' rule, what are they doing? How do they pass from nested set relations to a correct Bayesian answer? B&S offer no concrete or testable proposal, and thus no reason to exclude an associative solution. Why can't the transparency of the nested set relations allow other associative processes to kick in? It is question-begging to assume that the associative system is a priori unable to solve the task.

Indeed, there are at least two arguments that support this alternative possibility. First, our sensitivity to nested sets relations might itself rest on System 1 (associative) processes. When we look at the Euler circles, we simply "see" that one set is included in the other (perhaps this is why they are so useful, because they recruit another System 1 process?). Second, it is not hard to conceive of an associative system that gives correct answers to the Medical Diagnosis problem. Such a system just needs to learn that the correct diagnosis covaries with the base rate as well as the test results. This could be acquired by a simple network model trained on numerous cases with varying base rates and test results. And a system (or person) that learned in this way could be described as implementing the correct Bayesian solution.

The dual-process framework in general makes a strong distinction between normative and non-normative behaviour. In so doing, it embraces everything and explains nothing. One simply cannot align the normative/non-normative and rule-based/associative distinctions. True, rule-based processes might often behave in accordance with a norm such as Bayes' theorem, and associative systems non-normatively (as in the example from Gluck & Bower 1988); but, as argued above, it is also possible for rule-based processes to behave irrationally (think of someone explicitly using an incorrect rule), and for associative systems to behave normatively (backpropagation networks are, after all, optimal pattern classifiers).

Moreover, we know that without additional constraints, each type of process can be enormously powerful. Imagine a situation in which patients with symptom A and B have disease 1, while those with symptoms A and C have disease 2, with the former being more numerous than the latter (i.e., the base-rate of disease 1 is greater). Now consider what inference to make for a new patient with only symptom A and another with symptoms B and C. Both cases are ambiguous, but if choice takes account of base-rate information, then disease 1 will be diagnosed in both cases. In fact, people reliably go counter to the base rate for the BC conjunction (hence the "inverse base-rate effect"), choosing disease 2, whereas they choose disease 1 for symptom A (Medin & Edelson 1988; Johansen et al., in press). Thus, in one and the same situation, we see both

usage and counter-usage of base-rate information. But strikingly, these simultaneous patterns of behaviour have been explained both in rule-based systems (Juslin et al. 2001) and in associative ones (Kruschke 2001), emphasizing the inappropriateness of linking types of behaviour (normative, non-normative) to different processing "systems" (rule-based or associative).

The crux of B&S's argument, that a dual-process framework explains people's performance on probability problems, is unconvincing both theoretically and empirically. This is not to dismiss their critique of the frequentist program, but to highlight the need for finer-grained analyses. A crude dichotomy between the associative-system and the rule-based system does not capture the subtleties of human inference.

Ordinary people do not ignore base rates

DOI: 10.1017/S0140525X07001811

Donald Laming

University of Cambridge, Department of Experimental Psychology, Downing Street, Cambridge, CB2 3EB, United Kingdom.

drjl@hermes.cam.ac.uk

Abstract: Human responses to probabilities can be studied through gambling and through experiments presenting biased sequences of stimuli. In both cases, participants are sensitive to base rates. They adjust automatically to changes in base rate; such adjustment is incompatible with conformity to Bayes' Theorem. "Base-rate neglect" is therefore specific to the exercises in mental arithmetic reviewed in the target article.

Text removed due to third party copyright