

2007), statistically sophisticated subjects who solved the MD problem (66% of 35 subjects) were not able to solve the computationally less complex Linda problem (only 14% of the subjects did not commit the conjunction fallacy). So, the conceptual distinction between two reasoning systems, which explains the biases recurring in System 1 and the normative performance as the result of the activation of System 2, gives rise to some doubts.

According to us, the ability of statistically sophisticated subjects to grasp the informativeness of the data and the aim of the task in Bayesian tasks is a pragmatic ability. Also, when subjects don't give the logical-normative solution to the Linda problem, they are again considering the informativeness of the data, which, in this instance, hinders the intent of the experimenter (concerning the inclusion-class rule), because of a misleading contextualization of the task.

We could further speculate that, instead of having an ability for decontextualizing the task (Stanovich & West 2000), those gifted subjects who give the normative solution to the Linda problem (14%) would have a high ability to understand which context the experimenter intended, thereby revealing an interactional intelligence.

NOTE

1. Except when the number of men and women is the same.

Nested sets theory, full stop: Explaining performance on Bayesian inference tasks without dual-systems assumptions¹

DOI: 10.1017/S0140525X07001835

David R. Mandel

Defence Research and Development Canada (Toronto), Toronto, ON M3M 3B9, Canada.

david.mandel@drdc-rddc.gc.ca

<http://mandel.socialpsychology.org/>

Abstract: Consistent with Barbey & Sloman (B&S), it is proposed that performance on Bayesian inference tasks is well explained by nested sets theory (NST). However, contrary to those authors' view, it is proposed that NST does better by dispelling with dual-systems assumptions. This article examines why, and sketches out a series of NST's core principles, which were not previously defined.

A voluminous literature documents the many shortcomings people exhibit in judging probability. Barbey & Sloman (B&S) focus on a subset of this research that explores people's abilities to aggregate statistical information in order to judge posterior probabilities of the form $P(H|D)$, where D and H represent data and a focal hypothesis, respectively. Some of this literature indicates that people neglect base rates, although some of the findings are consistent with other judgment errors, such as the *inverse fallacy* (Koehler 1996), which involves confusing $P(H|D)$ with $P(D|H)$. For instance, Villejoubert and Mandel (2002) observed that bias (i.e., systematic inaccuracy) and incoherence (i.e., nonadditivity) in posterior probability judgments was well explained by the inverse fallacy, even though base-rate neglect could not account for the observed performance decrements. Thus, there is some question regarding exactly how much of what has been called base-rate neglect is in fact base-rate neglect. A safer claim is that performance on such Bayesian inference tasks is often suboptimal and much of the error observed is systematic.

B&S challenge a set of theoretical positions oriented around the core notion that humans are better at judging probabilities when the information they are provided with is in the form of natural frequencies (Gigerenzer & Hoffrage 1995). They argue – convincingly, I believe – that variation in frequency versus probability formats neither explains away performance errors, nor

does it account for errors as well as variation in the transparency of the nested set structure of an inference task. I shall not repeat their arguments here. Rather, my aim is, first, to sketch out some key propositions of nested sets theory (NST), which have yet to be described as a series of interlocking principles. Second, I will argue that NST would be on even firmer theoretical ground if the dual-systems assumptions that currently pervade B&S's version of it were jettisoned.

At its core, NST consists of a few simple propositions: First, performance on a range of reasoning tasks can be improved by making the partitions between relevant sets of events more transparent. I call this the *representation principle*. Second, because many reasoning tasks, such as posterior probability judgment, involve *nested* set relations, transparency often entails making those relations clear as well. I call this the *relational principle*. Third, holding transparency constant, nested set representations that minimize computational complexity will optimize performance. I call this the *complexity principle*. Fourth, the manner in which task queries are framed will affect performance by varying the degree to which default or otherwise salient representations minimize task complexity. In effect, this is the flip side of the complexity principle, and I call it the *framing principle*. Fifth, improvements in the clarity of nested set representations can be brought about through different modalities of expression (e.g., verbal description vs. visual representation). I call this the *multi-modal principle*. Sixth, within a given modality, there are multiple ways to improve the clarity of representation. I call this the *equivinality principle*. This list is almost certainly incomplete, yet it provides a starting point for developing a more explicit exposition of NST, which up until now has been more of an assemblage of hypotheses, empirical findings, and rebuttals to theorists proposing some form of the “frequentist mind” perspective. In the future, attempts to develop NST could link up with other recent attempts to develop a comprehensive theory of the representational processes in probability judgment (e.g., Johnson-Laird et al. 1999; Mandel, in press).

Although NST is not intrinsically a dual-systems theory (DST), B&S have tried to make it “DST-compatible.” This is unfortunate for two main reasons. First, although DSTs are in vogue (for an overview, see Stanovich & West 2000) – perhaps because they offer a type of Aristotelian explanation long favored by psychologists (Lewin 1931) – they are not particularly coherent theoretical frameworks. Rather, they provide a rough categorization of the processes that guide reasoning and that influence performance through an effort-accuracy tradeoff. The second reason for preferring “pure NST” to an NST-DST hybrid is that the former is not only more parsimonious, it actually offers a better explanatory account. According to the hybrid theory, when the nested set structure of a reasoning task is unclear, people have difficulty applying the rigorous rule-based system (also called “System 2”) and fall back on the more error-prone associative system (also called “System 1”). However, B&S say little about how judgment biases arise from those associative processes, or how system switching may occur.

Pure NST does not preclude the idea that impoverished representations of nested set relations can shift judgment towards a greater reliance on associative reasoning processes, but nor does it depend on that idea either. A viable alternative explanation is that impoverished representations lead to performance decrements because they increase one's chances of failing to access the correct solution to a problem. This does not necessarily mean that they switch to associative processes. It may simply mean that they fail to apply the correct principle or that they select the wrong information aggregation rule. Consider the inverse fallacy: It seems more likely that the error stems from a failure to understand how to combine $P(D | \neg H)$ with $P(D | H)$ (and with the base rates of H and $\neg H$ where they are unequal) than that it follows from use of associative processes.

Improving the representational quality of nested sets may also influence rule-based processes by simplifying the computations

required to implement a normative information aggregation strategy. Indeed, as B&S indicate, the performance decrements on Bayesian judgment tasks that Girotto and Gonzalez (2001) observed when participants were presented with “defective” (but nevertheless transparent) nested sets, appear to be attributable to the fact that such representations require at least one additional computational (subtraction) step. That computation itself may not be difficult to perform, but if it is missed the participant’s judgment will surely be wrong.

In short, the types of errors that arise from impoverished representations of nested set relations are generally consistent with a rule-based system. NST should remain pure and single, unencumbered by a marriage to dual-systems assumptions.

NOTE

1. The author of this commentary carried out this research on behalf of the Government of Canada, and as such the copyright of the commentary belongs to the Canadian Crown and is not subject to copyright within the United States.

Naturally nested, but why dual process?

DOI: 10.1017/S0140525X07001847

Ben Newell and Brett Hayes

School of Psychology, University of New South Wales, Sydney 2052, Australia.

ben.newell@unsw.edu.au b.hayes@unsw.edu.au

<http://www2.psy.unsw.edu.au/Users/BNewell>

Abstract: The article by Barbey & Sloman (B&S) provides a valuable framework for integrating research on base-rate neglect and respect. The theoretical arguments and data supporting the nested set model are persuasive. But we found the dual-process account to be under-specified and less compelling. Our concerns are based on (a) inconsistencies within the literature cited by B&S, and (b) studies of base-rate neglect in categorization.

Why so low? A striking feature of the data reviewed by Barbey & Sloman (B&S) is that the percentage of participants achieving the correct answer in the Medical Diagnosis problem rarely exceeds 50% (e.g., Table 3 of the target article). Thus, whether it is presenting information as natural frequencies or making nested set relations apparent that leads to improvements, overall the levels of performance remain remarkably low.

Potential reasons for this low level of overall performance are not discussed adequately in the target article. Although acknowledging in section 2.2 that “wide variability in the size of the effects makes it clear that in no sense do natural frequencies eliminate base-rate neglect” (para. 2), B&S fail to apply the same standard to their own proposal that “set-relation inducing formats” (be they natural frequencies or otherwise) facilitate a shift to a qualitatively different system of reasoning. The clear message of the article is that by presenting information appropriately, participants can “overcome their natural associative tendencies” (sect. 4, para. 3) and employ a reasoning system which applies rules to solve problems. Why does this system remain inaccessible for half of the participants in the studies reviewed? Is the rule system engaged, but the wrong rules are applied (e.g., Brase et al. 2006, Experiment 1)? Or do these participants remain oblivious to the nested sets relations and persevere with “inferior” associative strategies?

B&S cite evidence from studies of syllogistic reasoning, deductive reasoning, and other types of probability judgment in support of their contention that nested set improvements are domain-general. In these other tasks, however, the improvements are considerably more dramatic than in the base-rate studies (e.g., Newstead [1989] found a reduction in errors from 90% to 5% for syllogisms with Euler circle representations). The contrast between these large improvements in other

domains and the modest ones in the base-rate neglect problems sits uncomfortably in a dual-process framework. Why doesn’t the rule-based system overcome associative tendencies in similar ways across different tasks? In essence, the issue is one of specification – one needs to be able to specify what aspects of a problem make it amenable to being solved by a particular system for the notion of dual systems to have explanatory value. Why not simply appeal to “difficulty” and create a taxonomy or continuum of tasks differentially affected by various manipulations (diagrams, incentives, numerical format). Such a framework would not require recourse to dual processes or the vague rules of operation and conditions for transition that duality entails.

Incentives to “shift” systems? B&S report evidence from a study by Brase et al. (2006) in which it was found that monetary incentives improved performance on a base-rate problem. These data might be useful in gaining a clearer understanding of the factors that induce a shift between systems. It is difficult to make clear predictions, but under one interpretation of the dual-process position, one might expect incentives to have a larger effect in problems for which the set relations are apparent. The idea being that when the representation of information prevents (the majority of) people from engaging the rule-based system (e.g., when probabilities are used), no amount of incentive will help – most people simply won’t “get it.” A simple test of this hypothesis would be to compare probability and frequency incentive conditions. Brase et al. did not do this, comparing instead natural frequency conditions with and without additional pictorial representations. One would assume that the pictorial representations enhance nested set relations (target article, sect. 2.5) and increase the likelihood of shift to the rule-based system; hence, incentives would be more effective in the pictorial condition than condition. Brase et al.’s findings were telling: there was a main effect of incentives but this did not interact with the presence or absence of the picture; and indeed there was no main effect of providing a pictorial representation.

Two processes or two kinds of experience? In evaluating the B&S account we believe that it is useful to consider some of the lessons learned from the study of base-rate respect and neglect in category learning. In these studies people learn to discriminate between the exemplars of categories that differ in their frequency of presentation. The question is whether this base-rate information is used appropriately in subsequent categorization decisions, with features from more common categories being given greater weight. The results have been mixed. People can use base-rate information adaptively (Kruschke 1996), ignore it (Gluck & Bower 1988), or show an inverse base-rate effect, giving more weight to features from less frequent categories (Medin & Edelson 1988). Note that the issue of information format does not arise here, as all learning involves assessments of feature frequency. Critically, one does not need to invoke multiple processes to explain these results. Kruschke (1996) has shown that sensitivity and insensitivity to category base-rates can be predicted by a unitary learning model that takes account of the order in which different categories are learned, and allows for shifts of attention to critical features. In brief, people only neglect category-base rates when their attention is drawn to highly distinctive features in the less frequent category. The moral here is that before we resort to dual-process explanations of base-rate respect and neglect, we should first consider explanations based on the way that general learning mechanisms interact with given data structures.

Conclusion. B&S provide a very useful overview of the base-rate-neglect literature and provide convincing arguments for