principle, that $P(q|p)$ remains the same when computed on an enlarged space. This is the only way in which one can guarantee that enlargements of the probability space in the limit lead to a coherent probability distribution – the starting point of Bayesian rationality.

(2) An orthodox Bayesian alternative would be a construction in which the probability spaces remain the same (namely, the universal space based on all possible propositions), but the probability distributions change. In our toy world, the probability space is in both cases $\{p, q, r\}$, but one could assume that the probability distribution first assigns probability *0* to *not-r*, and, upon becoming aware of the second conditional "if *r* then *q*," a nonzero probability. The trouble with such a suggestion is that from a Bayesian point of view, the transition from the a priori probability $P(not\text{-}r)=0$ to the a posteriori $P(not\text{-}r) > 0$ is not allowed, because this cannot be achieved via (BaCo): conditionalizing on more evidence cannot make a null probability positive. One thus needs an additional rationality principle (beyond [BaCo]) governing such transitions. In the absence of such a principle, one has to assume that the probabilities of all non-salient exceptions (such as *not-r*) are initially very small but nonzero. This increases the computational complexity of probabilistic reasoning enormously: One requires massive storage and intricate computations to maintain consistency of the probability assignment.

These considerations show that in order to account for the data on the suppression task any probabilistic model needs to be supplemented with a theory about nonmonotonic and non-Bayesian, but still somehow rational, changes in degrees of belief. One may then question whether a probabilistic model is necessary at all; Stenning and van Lambalgen (2005; 2008a) provide a model cast entirely in terms of nonmonotonic logic.

The dynamics of development: Challenges for Bayesian rationality

doi:10.1017/S0140525X09000491

Nils Straubinger, Edward T. Cokely, and Jeffrey R. Stevens

Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Berlin, 14195, Germany.

straubinger@mpib-berlin.mpg.de

cokely@mpib-berlin.mpg.de

jstevens@mpib-berlin.mpg.de

http://www-abc.mpib-berlin.mpg.de/users/jstevens/

Abstract: Oaksford & Chater (O&C) focus on patterns of typical adult reasoning from a probabilistic perspective. We discuss implications of extending the probabilistic approach to lifespan development, considering the role of working memory, strategy use, and expertise. Explaining variations in human reasoning poses a challenge to Bayesian rational analysis, as it requires integrating knowledge about cognitive processes.

Bayesian rationality highlights the remarkable successes rather than failures of human reasoning by recasting seemingly erroneous reasoning in logical tasks using a probabilistic approach. However, in their book *Bayesian Rationality* (Oaksford & Chater 2007, henceforth BR), Oaksford & Chater (O&C) draw a rather static picture of human reasoning by focusing on typical patterns of responses from adults. We propose a more dynamic perspective, which considers that reasoning systematically varies within individuals over the lifespan (Howe & Rabinowitz 1996; Markovits & Barrouillet 2002) and between individuals with different levels of knowledge and expertise (Ericsson et al. 2006). Although O&C acknowledge the importance of considering reasoning data on individual differences (*BR*, p. 288) and on information processing capacities (p. 290), they do not adequately account for how variation influences a Bayesian rational analysis of reasoning. Anderson (1991a) and others have pointed out that perhaps the major potential limitation, the "Achilles heel," of rational analysis would be computational constraints that are too complex or arbitrary. We argue that our understanding of the mechanisms of change in reasoning can help us specify computational limitations for probabilistic modeling and assess whether a single model can capture the complexities of reasoning.

Many important aspects of cognition change over the lifespan, and reasoning is no exception (Baltes et al. 1999). According to Piaget, both logical reasoning and probabilistic reasoning emerge from adolescence to young adulthood at the highest stage of cognitive development (Piaget & Inhelder 1975). Subsequent research, however, has qualified these findings, showing that younger children understand aspects of such reasoning (Falk & Wilkening 1998; Galotti et al. 1997). Furthermore, reasoning continues to develop during adulthood with performance in specific domains increasing as individuals gain reasoning knowledge and expertise (Ericsson & Lehmann 1996; Sternberg 1999). Yet, overall across the adult lifespan, abstract reasoning (measured by intelligence tests) declines with age (Verhaeghen & Salthouse 1997). Thus, reasoning is a dynamic aspect of cognition that varies with age and experience and results from the interplay of biological processes (e.g., brain maturation) and enculturation (e.g., education) (Baltes et al. 1999).

A developmental perspective may inform Bayesian rational analysis by specifying computational limitations of the cognitive system. An important limitation faced by the human cognitive system is working memory capacity – a key determinant of reasoning performance (Kyllonen & Christal 1990). Like other cognitive capacities, working memory systematically changes across the lifespan by steadily increasing during childhood (Conlin et al. 2005) and declining across adulthood (Verhaeghen & Salthouse 1997). Working memory, therefore, poses a dynamic constraint on the rational analysis of reasoning.

Although O&C are currently silent on the role of developmental changes in working memory and reasoning, they do note that individuals with higher working memory capacities tend to exhibit more logical reasoning. To illustrate, in the Wason selection task, a subgroup of individuals (ca. 10%) consistently chooses the logically correct combination of cards, indicating that although most seem to adopt a probabilistic model, others clearly do not. O&C suggest that this variation in behavior primarily reflects deliberative strategy use and educational (training) differences, which are "not indicative of individual differences in the nature of the fundamental principles of human reasoning" (*BR*, p. 288). This claim seems problematic given what we know about the interplay between strategy use, training, and basic cognitive mechanisms. Of course, cognitive capacities can constrain the strategies that people use; however, specific strategy use and training may shape the basic cognitive mechanisms, as well. Differences in memory strategies (e.g., rehearsal, chunking) can also alter basic mechanisms of working memory capacity and its relationship to cognitive performance (Cokely et al. 2006). In addition, both extensive practice with specific strategies and the acquisition of knowledge and expertise dramatically expand working memory (Ericsson & Kintsch 1995). Indeed, as training changes deliberative strategies to automatic processes, the cortex can undergo functional neuroanatomical reorganization (Dick et al. 2006). Thus, it is possible that deliberative strategy use and training may influence reasoning precisely *because* they alter underlying cognitive mechanisms such as working memory. Given the complex relationship between strategies, training, and cognitive mechanisms, it seems premature to dismiss individual differences in strategy use as not fundamental to reasoning. A comprehensive model of human reasoning must account for these differences.

Variation in human reasoning has proven difficult to capture for probabilistic models (Shultz 2007), although recent research has made some progress applying probabilistic models to individual differences (e.g., category learning; Navarro et al. 2006) and cognitive development (e.g., causal reasoning; Sobel et al. 2004). This work represents a step in the right direction; however, we expect that no single model can predict reasoning performance equally well across age groups and levels of experience. Indeed, systematic variations in peoples' behavior suggest that several different models (or modifications of a given model) may be required to explain developing behavior (Shultz 2007). Nevertheless, investigating differences between the models across age groups and skill levels may help us to understand more exactly "what differs" between and "what develops" within individuals.

In closing, we must emphasize O&C's comment that probabilistic models are often only functional level theories that should not be confused with algorithmic level theories (process models). Brighton and Gigerenzer (2008) have pointed out in their discussion of the limits of Bayesian models of cognition that the question of why the human mind does what it does (functional level) cannot be separated from the question of how the human mind does it (algorithmic level). Therefore, it is crucial that future Bayesian rational analyses specify how exactly their functional level models constrain theorizing about cognitive processes. This issue is especially relevant as the data connecting development, expertise, working memory, and reasoning imply that multiple strategies (and therefore processes) are at play. Though Bayesian rationality seems to provide a functional level account of prototypical adult reasoning, the development of cognitive capacities and expertise remains underappreciated.

ACKNOWLEDGMENTS
Preparation of this commentary was supported in part by a stipend from the International Max Planck Research School LIFE to NS. We are grateful to Henrik Olsson for comments on a previous draft.

How do individuals reason in the Wason card selection task?

doi:10.1017/S0140525X09000508

Eric-Jan Wagenmakers
Department of Psychological Methods, University of Amsterdam, Amsterdam 1018 WB, The Netherlands.
EJ.Wagenmakers@gmail.com
http://users.fmg.uva.nl/ewagenmakers/

Abstract: The probabilistic approach to human reasoning is exemplified by the information gain model for the Wason card selection task. Although the model is elegant and original, several key aspects of the model warrant further discussion, particularly those concerning the scope of the task and the choice process of individuals.

In the book *Bayesian Rationality* (Oaksford & Chater 2007, henceforth *BR*), Oaksford & Chater (O&C) present a summary and a synthesis of their work on human reasoning. The authors argue that formal logic and deduction do not explain how people reason in everyday situations. The deficiencies of the most simple forms of logic are obvious when one considers that they may assign "true" to absurd statements such as "if the moon is blue, than cows eat fish" (*BR*, p. 70). More importantly, the authors propose that, in contrast to formal logic, probability calculus does provide the right tools for an analysis of human reasoning. Thus, the authors argue that people solve deductive tasks by inductive methods. From this perspective, human reasoning can be characterized as Bayesian or rational.

Consider the Wason card selection task discussed in Chapter 6. Participants are confronted with four cards, showing an A, a K, a 2, and a 7. Participants are told that each card has a number on one side and a letter on the other. They are given a rule, "if there is an A on one side, then there is a 2 on the other side," and subsequently, have to select those cards that need to be turned over to assess whether the rule holds true or not. A moment's thought reveals that the cards that need to be turned over are the A card and the 7 card. Yet, the majority of participants do not choose the 7 card, but tend to choose the 2 card instead.

O&C propose an elegant Bayesian model – the information gain model – to account for people's performance in the Wason task. According to the model, people select the cards that reduce their expected uncertainty the most. Specific assumptions about the rarity of the information on the cards lead to the conclusion that selection of the 2 card might be rational after all.

The information gain model has been subjected to intense scrutiny (e.g., Oberauer et al. 1999). For non-experts, the details of this discussion are somewhat difficult to follow. A useful guideline is that a model should only be abandoned when it can be replaced with something better. And – criticisms raised against the information gain model notwithstanding – I have not come across a model that does a better job explaining how people make their card selections.

Despite its simplicity and elegance, some important details of the information gain model were not clear to me. First, O&C argue, on page 210, that their account only holds if participants regard the cards as a sample from a larger population. Perhaps the authors could spell out this argument in a bit more detail. Taking probability as a reflection of degree of belief, I did not immediately see what calculations are in need of adjustment. Second, the authors mention that participants who realize that the cards are *not* sampled from a larger population would always choose the A card and the 7 card. I do not know whether this prediction has been tested empirically, but I find it only slightly more plausible than cows eating fish. Note that in the Wason task a substantial proportion of participants do not even select the A card.

Another issue that warrants closer examination is the way the model's predictions relate to the data. In the information gain model, each card reduces the expected uncertainty to some extent. Why then does an individual participant not select all four cards, but generally only selects one or two? In other words, it was unclear to me how the model, from a consideration of expected uncertainty reduction, can predict card selections for an individual participant.

A fourth point concerns the role of individual differences. As the authors discuss on page 211, a subgroup of undergraduate students with high intelligence (about 10%) do select the A card and the 7 card. This result strengthened my initial belief that a motivated, intelligent person would always choose the A and 7 cards, when given sufficient time. In the spirit of falsification, I then tested this assumption on a colleague, who of course immediately selected the A and 2 cards. Perhaps she was not sufficiently motivated to think the problem through carefully – would incentives of time or money increase the selection of the 7 card?

O&C are to be admired for their principled approach to quantitative modeling, and for their courage to take on the unassailable dogma of human irrationality. It is unfortunate that much of the material in the book was already available elsewhere (e.g., Oaksford & Chater 2001; 2003b); therefore, it was not entirely clear to me what the book adds to our current knowledge base.

One final comment. It strikes me as paradoxical that researchers who argue for a coherent, rational approach to human reasoning then proceed to apply an incoherent, irrational approach to the statistical analysis of their experimental data. Throughout the book, the authors renounce Popper's stance on the importance of falsification, arguing that this is not how science works, nor how people reason. But then, in the very same work, the authors measure the validity of their models by means of *p*-values, and include statements such as "the model could not be rejected." Why?