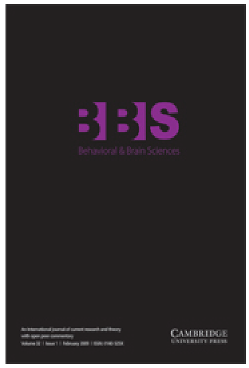

The uncertain reasoner: Bayes, logic, and rationality

Mike Oaksford and Nick Chater

Behavioral and Brain Sciences / Volume 32 / Issue 01 / February 2009, pp 105 - 120
DOI: 10.1017/S0140525X0900051X, Published online: 12 February 2009

Authors' Response

The uncertain reasoner: Bayes, logic, and rationality

doi:10.1017/S0140525X0900051X

Mike Oaksford[a] and Nick Chater[b]

[a]*School of Psychology, Birkbeck College London, London, WC1E 7HX, United Kingdom;* [b]*Division of Psychology and Language Sciences & ESRC Centre for Economic Learning and Social Evolution, University College London, London, WC1E 6BT, United Kingdom.*

mike.oaksford@bbk.ac.uk
n.chater@ucl.ac.uk
www.bbk.ac.uk/psyc/staff/academic/moaksford
www.psychol.ucl.ac.uk/people/profiles/chater_nick.htm

Abstract: Human cognition requires coping with a complex and uncertain world. This suggests that dealing with uncertainty may be the central challenge for human reasoning. In *Bayesian Rationality* we argue that probability theory, the calculus of uncertainty, is the right framework in which to understand everyday reasoning. We also argue that probability theory explains behavior, even on experimental tasks that have been designed to probe people's logical reasoning abilities. Most commentators agree on the centrality of uncertainty; some suggest that there is a residual role for logic in understanding reasoning; and others put forward alternative formalisms for uncertain reasoning, or raise specific technical, methodological, or empirical challenges. In responding to these points, we aim to clarify the scope and limits of probability and logic in cognitive science; explore the meaning of the "rational" explanation of cognition; and re-evaluate the empirical case for Bayesian rationality.

R1. Introduction

Bayesian Rationality (Oaksford & Chater 2007, henceforth *BR*) proposed that human reasoning should be understood in probabilistic, not logical, terms. In Part I, we discussed arguments from the philosophy of science, artificial intelligence, and cognitive psychology, which indicate that the vast majority of cognitive problems (outside mathematics) involve uncertain, rather than deductively certain, reasoning. Moreover, we argued that probability theory (the calculus for uncertain reasoning) is a more plausible framework than logic (the calculus for certain reasoning) for modeling both cognition in general, and commonsense reasoning in particular. In Part II, we considered a strong test of this approach, asking whether the probabilistic framework can capture human reasoning performance even on paradigmatically "logical" tasks, such as syllogistic reasoning or conditional inference.

The structure of this response is as follows. In section R2, we reflect on the ubiquity of uncertainty and address the theoretical attempts to preserve logic as a separate and core reasoning process. In section R3, we compare and evaluate Bayesian and logic-based approaches to human reasoning about uncertainty. Section R4 focuses on the methodology of rational analysis (Anderson 1990; 1991a; Oaksford & Chater 1998b) and its relationship to more traditional algorithmic and neuroscientific approaches. Section R5 discusses a variety of specific issues in the empirical data from the psychology of reasoning, and the modeling of that data. Finally, section R6 concludes the case for a "Bayesian turn" in the brain and cognitive sciences in general, and for the understanding of human reasoning in particular.

R2. The ubiquity of uncertainty: Distinctions that might preserve logic

Many commentators suggest ways to preserve a role for logic as a separate and core component in an account of human reasoning, despite the challenge provided by uncertainty (**Allott & Uchida**, **Evans**, **Politzer & Bonnefon**). We argue that logic does have an important role in modeling cognition; but we argue against the existence of cognitive processes dedicated to logical reasoning.

R2.1. Rationality 1 versus Rationality 2

Evans suggests that a distinction should be drawn between two types of rationality (Evans & Over 1996a). Rationality 1 relates to implicit, possibly associative, processes, operating over world knowledge, which Evans also terms "ecological rationality." This type of rationality arises from System 1 in Evans and Over's (2004) Dual Process Theory (see also Evans & Frankish, in press; Sloman 1996; Stanovich & West 2000). Rationality 2 involves explicitly following normative rules, and is the type of rationality achieved by Evans and Over's (2004) System 2. System 2 processes are logical, rule-governed, and conscious. Moreover, Evans has argued for a crucial asymmetry between the systems. It requires cognitive effort to ignore System 1, and to use System 2 for logical inference: that is, to infer only what follows from the structure of the given premises.

The fundamental problem with this Dual Process view is that these two systems must interact – and if the systems obey fundamentally different principles, it is not clear how this is possible. Consider the familiar example of inferring that *Tweety flies* from the general claim that *birds fly* and the fact that *Tweety is a bird*. On the Dual Process view, this inference could be drawn logically from the premises given by System 2, from the assumption that *birds fly* is a true universal generalization; System 1, by contrast, might tentatively draw this conclusion by defeasible, associative processes, drawing on general knowledge. But a lack of synchrony between the two systems, presumed to operate by different rational standards, threatens to cause inferential chaos. Consider, for example, what happens if we consider the possibility that *Tweety is an ostrich*. If System 2 works according to logical principles, the clash of two rules threatens contradiction: we know that birds fly, but that ostriches do not. To escape contradiction, one of the premises must be rejected: most naturally, *birds fly* will be rejected as false. But we now have two unpalatable possibilities. On the one hand, suppose that this retraction is *not* transferred to general knowledge and hence is not assimilated by System 1. Then the two systems will have contradictory beliefs (moreover, if System 2 reasoning cannot modify general knowledge, its purpose seems unclear). On the other hand, if *birds fly* is retracted from world knowledge, along with other defeasible generalizations, then almost all of general knowledge will be stripped away (as *BR* notes, generalizations outside mathematics are typically defeasible), leading System 1 into inferential paralysis.

The centrality of logic for a putative System 2 is also brought into doubt by considering that one of its main functions is to consciously propose and evaluate *arguments*. Yet, argumentation, that is, the attempt to persuade oneself or others of a controversial proposition, is uniformly agreed not to be a matter of formal logic (Walton 1989), although aspects of argumentation may naturally be modeled using probability theory (Hahn & Oaksford 2007). Thus, perhaps the core human activity for which a logic-based System 2 is invoked may, ironically, be better explained in probabilistic terms.

People can, of course, be trained to ignore some aspects of linguistic input, and concentrate on others – for example, in the extreme, they can learn to translate natural language statements into predicate logic (ignoring further aspects of their content) and employ logical methods to determine what follows. But, for the psychology of reasoning, this observation is no more significant than the fact that people can learn the rules of chess and ignore most of the visual features of the pieces, the board, or indeed, their surroundings. Conscious application of logical principles is a learned skill built on top of non-logical machinery (and, indeed, is highly effortful, even for logicians); it does not involve, we suggest, tapping in to some underlying logical "engine."

It is this conscious application of logical concepts (and related notions from mathematics, philosophy, and computer science) that underpins, we suggest, the small but significant correlation between "logical" performance on some reasoning tasks (e.g., selecting the p and *not-q* cards, in Wason's selection task) and IQ (Stanovich & West 2000). Logical reasoning is a late and cognitively challenging cultural innovation, rather than a core component of our mental machinery.

Evans also expresses disappointment that we do not address individual differences (Stanovich 2008), which have been viewed as supporting a Dual Process account. But from the present perspective, individual differences concerning the application of learned logical rules are no different from individual differences in chess playing – that is, neither are directly relevant to the question of whether there are single or multiple reasoning systems. Indeed, we suggest that individual differences provide no stronger evidence that cognition involves core *logical* competence, than that cognition involves core *chess-playing* competence.

It may turn out, indeed, that there is no real incompatibility between Stanovich's account and ours. In particular, the distinction Stanovich draws between control processes and other autonomous systems is a distinction common to all theories of cognition (see Oaksford & Chater, in press). But as Kowalski's (1979) classic equation, "Algorithm = Logic + Control," reminds us, logic and control processes are very different (see, e.g., Anderson 1983). Hence, Stanovich may not really be committed to anything like **Evans**' logically competent System 2. (A further complication is that a distinction between processes of logic and control is now reflected in Evans [2007], who moots the possibility of a tri-process theory.)

R2.2. The split between semantics and pragmatics

Grice's (1975) theory of conversational implicature originally attempted to split off a "stripped down" logic-based natural language semantics, from the complex, knowledge-rich processes of pragmatic interpretation involved in inferring a speaker's intentions. In this way, he aimed to retain a logical core to semantics, despite apparently striking and ubiquitous clashes between the dictates of formal logic and people's intuitions about meaning and inference.

Within this type of framework, **Allott & Uchida** attempt to preserve the truth of potentially defeasible conditionals (*if it's a bird, then it flies*, or, as above, *birds fly*) despite the ready availability of counterexamples. They suggest that this conditional is true in one model, but not in the model that is considered when an additional premise giving a counterexample is added (e.g., when we consider the possibility that Tweety is an ostrich). But in classical logic, only an inference that holds in *all* models is deductively valid, by definition. Thus, accepting that this inference holds only in some models implies accepting that the inference is uncertain (contra, e.g., **O'Brien**). Indeed, in *BR*, we argue uncertainty is ubiquitous in human reasoning; outside mathematics, deductive reasoning, which *guarantees* the truth of a conclusion given the premises, is, to a first approximation, never observed.

Moreover, understanding reasoning involves working out pragmatic details about what default background assumptions are applicable in reasoning. Thus, for example, our accounts of specific reasoning phenomena, across conditional reasoning, the selection task, and syllogistic reasoning, involve default assumptions about the environment, for example, what is rare and what is common (cf. **McKenzie**; McKenzie et al. 2001) and when states are likely to be independent or conditionally independent. In this light, we agree with **Stenning & van Lambalgen**'s claim that "pure" Bayesian analysis, working from the premises alone, cannot capture suppression effects in conditional reasoning (see sect. R3.6) – we view this as illustrating the knowledge-rich character of reasoning, rather than challenging a Bayesian account.

The ubiquity of uncertain, knowledge-rich inference, argues for an alternative to invoking the semantics/pragmatics distinction to maintain a logical semantics for natural language: namely, that natural language semantics may be probabilistic "all the way down." Experiments in the psychology of reasoning, as reviewed in *BR*, find little support for the existence of a level of logic-based representation or inference. *BR* proposes a starting point for a probabilistic semantics: *If p then q* conditionals are assumed to express that the conditional probability $P(q|p)$ is high (following Adams 1975; 1998; Bennett 2003; and Edgington 1995, among others); the quantifiers *Some*, *Few*, *Most*, *All* are similarly assumed to express constraints on probabilities (e.g., *Some A are B* is rendered as $P(A, B) > 0$; *Most A are B* claims that $P(B|A)$ is high). Switching from a logical to a probabilistic semantics provides, we argue, a better fit with patterns of human reasoning. Of course, it remains possible that a logical core interpretation might be maintained – but it seems theoretically unparsimonious to do so (Edgington 1995).

A shift from a logical to a probabilistic semantics for aspects of natural language may also allow a more integrated account of semantics and pragmatics. Indeed, **McKenzie** (e.g., Sher & McKenzie 2006) has powerfully demonstrated the importance of pragmatic factors, even within a purely probabilistic framework (but see, Hilton et al. 2005). Nonetheless, the core insight of Grice's program remains: that splitting apart semantic factors (concerning meaning) and pragmatic factors (concerning inferences about speaker intentions) is a prerequisite for constructing a tractable semantic theory, whether that theory be based on logic (as **Allott & Uchida** argue) or probability (as *BR* proposes).

R2.3. Proof and uncertainty and structure and strength

Politzer & Bonnefon argue that a key element missing from a purely probabilistic account is how premises can be used to construct proofs to derive conclusions. Thus, they argue that the probabilistic account allows the *evaluation* of the strength of the relationship between premises and conclusion, but not how the conclusion is *generated* in the first place. Note, though, that both logic and probability are theories of the nature of inferential relationships between propositions (Harman 1986). Neither specify how reasoning should be carried out, let alone how interesting conclusions should be generated. Moreover, for both logic and probability, a range of algorithms have been developed which can both evaluate given conclusions, and generate new conclusions (e.g., logic programming and Bayesian networks). From both perspectives, any set of information potentially generates an infinite set of possible conclusions; so that an immediate question is: What counts as an *interesting* conclusion? A natural suggestion from the probabilistic point of view is that conclusions with a low prior probability are, other things being equal, more surprising and hence more interesting (as employed in the account of syllogistic reasoning described in *BR*), although interesting logic-based measures of semantic information content have also been proposed (Johnson-Laird 1983).

More generally, the probabilistic approach is just as able as logic-based approaches to serve as the basis for algorithmic models of thought. For example, Oaksford & Chater (in press) use a constraint satisfaction neural network implementation of the probabilistic approach. The links in the network captures the conditional and default assumptions about *structural* relations between variables (in the causal context, involving alternative causes and defeaters); and the *strength* of each link is captured by a weight. A similar distinction between structure and strength has been invoked in causal reasoning using Bayesian networks (Griffiths & Tenenbaum 2005) and applied in Hahn and Oaksford's (2007) probabilistic account of argumentation.

R3. Logic, probability, and the challenge of uncertain reasoning?

In this section, we consider whether, as some commentators suggest, we have mischaracterized the scope of logic or chosen the wrong alternative calculus in order to reason about uncertainty. We deal with logic and probability in turn.

R3.1. How are logic and probability related?

Pfeifer & Kleiter observe that probability theory already includes classical propositional logic as a special case. Thus, one way of understanding the approach outlined in *BR* is as enriching conventional logic to give an *inductive* logic – a system of logic that extends deduction to less-than-certain inferences (Hawthorn 2008). To a good approximation, modern inductive logic just *is* Bayesian probability (Chater et al., in press; Earman 1992), with some additional discussion of the measure of the confirmation relation (see later discussion of **Poletiek** and **Nelson**). Since Carnap (1950), this Bayesian inductive logic includes classical logic – if a statement has a probability of 1, then any logical consequence of that statement also has a probability of 1. Similarly, if a statement has an implication with a probability of 0, then that statement has a probability of 0 (note, however, that probability theory does not readily represent the *internal* structure of atomic propositions, and has no general theory of, for example, quantification or modality). The Bayesian inductive perspective is required not because classic logic is incorrect, but because, outside mathematics, it rarely, if ever, applies (Oaksford & Hahn 2007) – inferential relations between propositions are relentlessly uncertain (Jeffrey 1967).

R3.2. Is relevance relevant?

O'Brien proposes a different enrichment of logic, drawing on his important work with Braine on mental logics (Braine & O'Brien 1991), which aims to capture a notion of *relevance* between antecedent and consequent (i.e., so that conditionals such as *if 2 is odd, then the sky is purple* are no longer automatically true, just in virtue of the false antecedent). Thus, Braine and O'Brien's work aims to go beyond the material conditional, which *BR* ascribed to mental logic as a whole (e.g., Rips 1994).

Adding a condition of relevance, while potentially important, does not help deal with the problem of uncertain reasoning, however. Indeed, **O'Brien**'s account of conditionals is, instead, a strictly deductive version of the Ramsey test (like, e.g., Gärdenfors 1986) – conditionals are only asserted if the consequent, q, follows with certainty from the antecedent p (and background knowledge). Thus, Braine and O'Brien's (1991) logical interpretation of the conditional suffers the same fundamental problem as material implication: an inability to capture the fact that generalizations outside mathematics are inevitably uncertain.[1]

Moreover, despite Braine and O'Brien's intentions, their system does not seem to enforce relevance between antecedent and consequent, either. The introduction rule for *if p then q*, used by **O'Brien**, and described in Braine and O'Brien (1991), states that *if p then q* can be inferred if q follows from the supposition of p together with background knowledge, B. If we know p is false (i.e., background knowledge B implies *not-p*), then supposing p and B implies *p & not-p*, which is a contradiction, from which any conclusion follows – including q. So conditionals such as *if 2 is odd, then the sky is purple* can be asserted, after all. Similarly, any conditional whose conclusion is known to be true (i.e., B implies q) will automatically meet the condition that p & B implies q (because this is a monotonic logic – adding premises can never remove conclusions). Hence, conditionals such as *if the sky is purple, then 2 is even*, will also be asserted – again violating intuitions of relevance.

R3.3. Uncertain reasoning via nonmonotic logic?

Stenning & van Lambalgen argue that we misrepresent the scope of current logical methods, noting that a range of nonmonotonic logics, in which adding a premise may require withdrawing a previously held conclusion, might meet the challenge of uncertainty. As noted in *BR*, and elsewhere (e.g., Oaksford & Chater 1991; 2002), there are, however, fundamental problems for nonmonotonic logics in the crucial case where different "lines of argument" clash. Thus, *if it is sunny, John goes to the park*, and *it's sunny* appears to provide a powerful argument that *John goes to the park*. But adding the premise, *John is arrested by the police in a dawn raid*, together with background knowledge, appears to yield the conclusion that *John does not go to the park*.

From the perspective of classical logic, this situation is one of contradiction – and what is needed is a way of resolving which premise should be rejected. For example, one might claim that the conditional *if it's sunny, John goes to the park* is false, precisely because of the possibility of, among other things, arrest. But, as noted in section R2.1, it is then difficult to avoid the conclusion that all conditionals, outside mathematics, are false, because the possibility of counterexamples always exists. Reasoning from premises known to be false is not, of course, justified, whether in logic, or any other standard framework, and hence, the logical analysis of the original argument collapses.

The strategy of nonmonotonic logic attempts to solve this problem by treating the conditional as a *default* rule, which holds, other things being equal. Indeed, outside mathematics, almost all rules are default rules. Indeed, the implicit rule that allows us to infer that being arrested is incompatible with a trip to the park is itself a default rule, of course – for example, arrest may be extremely brief, or perhaps the police station is itself in the park. Thus, from this viewpoint, uncertain reasoning centrally involves resolving clashes *between* default rules. In *BR*, we argue that resolving such clashes is not typically possible by looking only at the structural features of arguments. Instead, it is crucial to differentiate stronger and weaker arguments, and degrees of confidence in the premises of those arguments. Logical methods provide no natural methods for expressing such matters of degree; but dealing with degrees of belief and strength of evidence is the primary business of probability theory.

R3.4. Is logic relevant to cognition?

Several commentators suggest that the powerful machinery of logic should not be jettisoned prematurely (**Allott & Uchida**, **De Neys**, **O'Brien**, **Politzer & Bonnefon**, **Stenning & van Lambalgen**). As we noted in section R3.1, probability theory (i.e., modern inductive logic) is a generalization of logic, allowing degrees of uncertainty. However, it is a generalization that is presently limited in scope. This is because how probability interacts with richer representations involving, for example, relations, quantification, possibility, deontic claims, tense and aspect, and so on, is yet to be worked out. *BR* has, as we have mentioned, some preliminary suggestions about the probabilistic representation of individual connectives (*if...then...*) and quantifiers (*Most*, *Few*, *Some*, etc.). But this is very far from a full probabilistic generalization of, for example, the predicate calculus, the workhorse of classical logic and natural language semantics. The formal challenges here are substantial. Nonetheless, much progress has been made, in a number of directions, in fusing together probabilistic and logical methods (e.g., see papers in Williamson & Gabbay 2003), thus advancing Carnap's (1950) program of building an inductive logic. **Pfeifer & Kleiter** apply logic in an interesting, but distinct, way: as providing a machinery for reasoning *about* probability, rather than using probability to generalize logic.

According to **De Neys**, concentrating on the computational level means that *BR* underplays the role of logic in human reasoning. De Neys argues that latency and brain imaging studies, investigating the mental processing involved in reasoning, rather than just the output of these processes, consistently reveal a role for logic. Yet all the cases that De Neys cites involve a conflict between belief and logic such that prior belief suggests one response, but logical reasoning from the given premises suggests another. However, the Bayesian approach can explain at the computational level why such conflicts might arise and therefore why inhibitory processes might need to be invoked (De Neys et al. 2008; Houdé et al. 2000). Oaksford and Hahn (2007) point out that probabilistic validity of an argument and its inductive strength can conflict. So, for example, Modus Ponens (MP) is probabilistically valid. However, if the probability of the conditional is low, then the inductive strength of the argument, that is, the probability of the conclusion given the premises, will also be low. The right computational level analysis may, therefore, remove the need to propose two special purpose cognitive systems operating according to different principles. This view seems consistent with the current state of imaging studies, which provide little evidence for a dedicated logical reasoning module (Goel 2007).

O'Brien describes Chrysippus' dog's ability to follow a scent down one path in a fork in the road, having eliminated the other as an application of the logical law of disjunction elimination – and hence, suggests that logic is cognitively ubiquitous. However, this logical law cannot uncritically be imported into a theory of canine cognition. For one thing, such patterns of behavior are at least as well modeled in probabilistic (Toussaint et al. 2006), as in logical, terms. Indeed, probabilistic methods are crucial in planning tasks in uncertain environments, which is, of course, the normal case, outside mathematically specified game-playing environments. In any case, just because a behavior can be *described* in logical or probabilistic terms does not directly imply that it is *governed* by logical or probabilistic processes. The issues here are complex (see the excellent introductory chapter to Hurley & Nudds 2006) and many possibilities would need to be ruled out before abandoning Lloyd Morgan's canon: that lower-level explanations of animal behavior should be preferred.

In short, we believe that cognitive science ignores logic at its peril – logic provides powerful and much needed tools, just as do other branches of mathematics. It does not, however, readily capture patterns of human reasoning, or, we suggest, cognition at large, unless generalized into a probabilistic form able directly to deal with uncertainty.

R3.5. Why probability rather than other numerical measures?

Danks & Eberhardt and **Politzer & Bonnefon** ask why we use probability, rather than other numerical measures of degrees of belief, such as confidence intervals, Dempster-Shafer belief functions (Dempster 1968; Shafer 1976), or fuzzy logic (Zadeh 1975). In *BR*, our primary motivation is practical: Bayesian probabilistic methods provide a natural way to capture human reasoning data; and more generally, Bayesian methods have swept through the brain and cognitive sciences, from understanding neural coding (Doya et al. 2007), through vision, motor control, learning, language processing, and categorization. Even within research on reasoning, Bayesian methods have proved central to understanding inductive inference (Griffiths & Tenenbaum 2005; Tenenbaum et al. 2007), causal reasoning (Sloman 2005; Tenenbaum & Griffiths 2001), and argumentation (e.g., Hahn & Oaksford 2007), as well as the primarily deductive reasoning problems considered in *BR*.[2] Moreover, probabilistic methods connect with rich literatures concerning computational inference methods (e.g., based on graphical models, Lauritzen & Spiegelhalter 1988; Pearl 1988), machine learning (e.g., Jacobs et al. 1991), and normative theories of reasoning about causality (Pearl 2000). Finally, probability also has deep relationships to other powerful concepts in the brain and cognitive sciences, including information theory (e.g., Blakemore et al. 1991) and simplicity, for example, as captured by Kolmogorov complexity theory (e.g., Chater 1996; Chater & Vitányi 2002). Thus, our focus on probability is primarily pragmatic rather than, for example, depending on a priori justifications.

Danks & Eberhardt focus, nonetheless, on justification, arguing that doubt can be cast on justifications such as the Dutch Book argument and long run convergence theorems. We see the project of rational analysis as a *user* of probability, on a par with the rest of science, for example, statistical mechanics, Bayesian image restoration, or economics. We only need to be as concerned about justification as these other endeavors. Danks & Eberhardt's worries are analogous to Berkeley's objections to Newton's infinitesimals: of considerable conceptual importance, but with little direct impact on the practical conduct of science. Nonetheless, probability is at least *better* justified than alternative formalisms for modeling uncertainty.

Politzer & Bonnefon and **Danks & Eberhardt** raise the possibility that the assumptions of the probabilistic approach may be too strong. We instead believe that they are, if anything, too weak; that is, they define minimal coherence conditions on beliefs, which need to be supplemented with richer formalisms, including, as noted in section R3.4, the ability to represent relations and quantification, and to represent and manipulate causal relations (e.g., Pearl 2000).

R3.6. Are we Bayesian enough?

Other commentators (**Over & Hajichristidis**, **Pfeifer & Kleiter**, **Stenning & van Lambalgen**) have the opposite concern: that *BR* is not Bayesian enough. Over & Hadjichristidis argue that in conditional inference, not only is the conditional premise (e.g., *if p then q*) uncertain, but so is the categorical premise, *p*. In *BR* (p. 121), we mention this general case (implying Jeffrey's rule [Jeffrey 1983]), but point out that this extra element of uncertainty appears unnecessary to capture the conditional reasoning data.

Stenning & van Lambalgen and **Pfeifer & Kleiter** also argue, in different ways, that we are insufficiently Bayesian. Stenning & van Lambalgen argue that our account of suppression effects is not Bayesian because coherent Bayesian revision of the probability space assumes "rigidity": that is, the conditional probability $P(q|p)$ remains the same if we learn the truth of a categorical premise: *p*, *q*, *not-p*, or *not-q* (and no other information). We agree. But this does *not* imply that $P(q|p)$ remains the same if we are *told* about that *p*, because pragmatic factors allow us to infer a great deal of additional information; and this information can legitimately change $P(q|p)$. It is this latter case that is relevant for reasoning with verbal materials. Thus, suppose I believe *if the key is turned, the car starts*; and I am told: "the car didn't start this morning." This would be a pragmatically pointless remark, if the key had not been turned. I therefore infer that the key was turned, and the car didn't start for some other reason. Thus, I revise down the probability of the relevant conditional P(*car starts*|*key turned*) dramatically. So the violation of rigidity, notably in this type of Modus Tollens (MT) inference, does not violate Bayesian precepts, but merely applies them to the pragmatics of utterances (see *BR*, pp. 126–128; Sobel 2004; Sober 2002).

Pfeifer & Kleiter suggest that inference can proceed locally and deductively in a mental probability logic. In such a logic, the precise probability of a conclusion cannot typically be deduced from the probabilities of the premises – but a probability *interval* can be. We adopted a similar approach to probabilistic validity for syllogisms where, according to our probabilistic semantics, quantifiers describe probability intervals. Nonetheless, in line with Stanovich and West's (2000) "fundamental computational bias," we believe that people spontaneously contextualize and elaborate verbal input, by adding information from world knowledge. Indeed, it takes substantial cognitive effort *not* to do this. Consequently, we think it unlikely that people reason deductively about probability intervals.

R3.7. Measuring confirmation

People are not merely passive observers. They can actively search for information to help test hypotheses, or to achieve specific goals. In *BR*, we outline "rational" accounts for both cases. Where people test between hypotheses, a natural objective is to search for data in order to maximize the expected amount of *information* that will be gained in the task (Shannon & Weaver 1949). This is "disinterested" inquiry. Where people gain information to help achieve specific goals, then a natural objective is to choose information to maximize expected utility (balancing costs of information search with the improved choices that may result from new information). This is "goal-directed" inquiry. In *BR*, we note that different variations of Wason's selection task are appropriately captured by versions of one or other model. In particular, we showed how attending to the goal-directed case avoids the postulation of specific machinery, such as "cheater-detection" modules (e.g., Cosmides 1989), to explain patterns of experimental data (e.g., *BR*, pp. 191–98).

Focusing on disinterested inquiry, **Nelson** notes that a wide range of normative and descriptive proposals for assessing the strength of information in a piece of data have been proposed. In testing these models against a wide range of psychological data (Nelson 2005), he finds that the information-theoretic measure implicit in our analysis stands up well against competitors, although it is not picked out uniquely by the empirical data.

Poletiek notes a further interesting link to philosophy of science, noting that Popper's measure of severity of test is equivalent to $P(e|H)/P(e)$, for data e and hypothesis H. And the logarithm of this quantity just *is* the amount of information carried by the evidence e about H – the quantity which we use in our model of disinterested inquiry in the selection task. This quantity is also used as a measure of the degree to which a theory is confirmed by the data in confirmation theory (Milne 1996). This is, as Poletiek notes, particularly interesting, given that Popper's measure of severity of test is part of a theoretical framework which aims to entirely avoid the notion of confirmation (see also Milne 1995). Thus, our account of the selection task could be recast, from a Popperian standpoint, as a rational analysis in which people attempt to choose data to provide the more severe possible tests for their hypotheses.

R4. Rational analysis, algorithmic processes, and neural implementation

BR is primarily concerned with the rational analysis of human reasoning (e.g., Anderson 1990; 1991a; Chater & Oaksford 2008a; Oaksford & Chater 1998b). In this section, we consider the role of rational analysis in the brain and cognitive science and whether this style of explanation is fundamentally flawed.

R4.1. The power of rational analysis

Hahn notes that the shift away from considerations of algorithms and representations, encouraged by rational analysis, has led to a substantial increase in explanatory power in cognitive science, in a number of domains. Where the underlying explanation for an aspect of cognition arises from the rational structure of the problem being solved, there focusing on specific algorithmic and neural mechanisms may be unhelpful. Therefore, building specific algorithmic models (e.g., connectionist networks) of a phenomenon may replicate the phenomenon of interest (by virtue of being an adaptive solution to the "rational" problem in hand), but may throw little light on *why* it occurs.

R4.2. Normativity and rational analysis

Evans and **Schroyens** are concerned about the normative aspect of rational analysis. Evans questions whether normativity is a proper part of a computational-level analysis of human reasoning, and by implication, cognition in general, and recommends a switch to an *ecological* notion of rationality. He suggests rationality should concern how well people are adapted to their environment, which may not require following the prescriptions of any normative theory of reasoning (cf. Gigerenzer & Goldstein 1996).

We suggest, however, that ecological rationality does not replace, but rather, complements normative rationality. Normative considerations are still required to explain *why* a particular algorithm works, given a particular environment; indeed, this is precisely the objective of rational analysis. Thus, for example, in arguing for the ecological rationality of various fast and frugal heuristics (Gigerenzer et al. 1999), Gigerenzer and colleagues appeal to a Bayesian analyses to explore the type of environmental structure for which their algorithms succeed (e.g., Martignon & Blackmond-Laskey 1999). Thus, rational analysis cannot be replaced by, but seeks to explain, ecological rationality.

Note, too, that rational analysis is goal-relative: it specifies how best to achieve a given goal, in a given environment, with given constraints (Anderson 1990; Oaksford & Chater 1998b). So, if your goal is to land a rocket on the moon, your guidance system ought to respect classical physics; if your goal is to avoid contradictions, you ought to reason according to standard logic; and if your goal is to avoid accepting bets that you are bound to lose, you ought to follow the rules of probability theory (see Ch. 2 of *BR*).

Ignoring the goal-relativity of rational analysis leads **Schroyens** to suggest that we have fallen into Moore's (1903) "naturalistic fallacy" in ethics: that we have attempted to derive an "ought" from an "is." Moore's concern is that no facts about human behavior, or the world, can justify an ethical theory. Ethics is concerned with non-relative notions of "ought": the aim is to establish universal principles of right behavior. But the goal-relativity of rational analysis makes it very different from the domain of ethics, because it is *conditional*. Rational analysis considers: *if* you have objective O, given an environment E, and constraints C, then the optimal action is A. Ethics, by contrast, considers whether O is a justifiable objective. And the nature of the solution to a well-specified optimization problem is itself firmly in the domain of facts.

Indeed, were **Schroyens**' concern valid, then its consequences would be alarming, sweeping away functional explanation in biology and rational choice explanation in economics. Yet in all cases, rational/optimality explanations are used to derive empirical predictions; and, as in any scientific enterprise, the assumptions of the rational/optimality accounts are adjusted, where appropriate, to give a better fit with empirical predictions. Specifically, empirical data lead to revision of *empirical* assumptions in the rationality/optimality analysis – the empirical data does *not* lead to a revision of the laws of logic, probability, or any other rational theory.

Khalil raises the opposite concern: that we use rational explanation too narrowly. He argues that the style of optimality explanation that we advocate applies just as well in the explanation of non-cognitive biological structures as it does to cognitive processes – he argues that, in the sense of rationality used in *BR*, stomachs are just as rational as cognitive mechanisms. This concern appears purely terminological; we reserve "rationality" for information processing systems. But rational analysis is, indeed, parallel to optimality explanation in biology (Chater et al. 2003).

R4.3. Relevance of the algorithmic level

McKenzie and **Griffiths** note, however, that advocating rational analysis does not make the challenges concerning algorithmic, and indeed neural, implementation, disappear. Moreover, the mapping between levels of explanation need not necessarily be straightforward, so that a successful probabilistic rational analysis of a cognitive task does not necessarily require that the cognitive system be carrying out probabilistic calculations – any more than the bird is carrying out aerodynamic calculations in growing a wing perfectly adapted for flight. Nonetheless, in many contexts, it *is* natural to see cognition as carrying out probabilistic calculations; and a prior rational analysis (or, in Marr's [1982] terms, *computational* level of explanation) is extremely valuable in clarifying what calculations need to be carried out. Without a "rational analysis" for arithmetic calculations (i.e., a mathematical theory of elementary arithmetic), understanding which algorithms might be used by a pocket calculator, let alone how those algorithms might be implemented in silicon, would be impossible. Griffiths outlines key challenges for creating an algorithmic-level theory of cognition, viewed from a Bayesian perspective; and this perspective dovetails nicely with work viewing neural machinery as carrying out Bayesian inference (e.g., Ma et al. 2006; Rao et al. 2002), which we consider briefly further on.

BR is largely focused on rational level explanation (Anderson 1990; 1991a). Indeed, following Marr (1982), we argued that understanding the rational solution to problems faced by the cognitive system crucially assists with explanation in terms of representations and algorithms, as stressed by **Hahn** and **Griffiths**. In *BR*, this is illustrated by our model of syllogistic reasoning, which proposes a set of "fast and frugal" heuristics (Gigerenzer & Goldstein 1996) for generating plausible conclusions, rooted in a Bayesian rational analysis (Chater & Oaksford 1999b). More recently, we have suggested methods for causal and conditional reasoning, based on "mental mechanisms" (Chater & Oaksford 2006; Ali et al., in press) directly building on rational and algorithmic models inspired by the literature on Bayesian networks (Glymour 2001; Pearl 1988; 2000). Moreover, an explicit algorithmic implementation of our probabilistic account of conditional inference has been constructed using a constraint satisfaction neural network (Oaksford & Chater, in press). Moreover, there is a significant movement in current cognitive science that focuses on developing and employing Bayesian machine learning techniques to model cognition at both the rational and algorithmic levels (e.g., Griffiths et al. 2007; Kemp & Tenenbaum 2008).

Evans' concern that we ignore the algorithmic level is therefore puzzling. He worries that *BR* recommends that one should "observe some behaviour, assume that it is rational, find a normative theory that deems it to be so, and then . . . nothing else, apparently." We assume that the ellipsis should, in Evans' view, be fleshed out with an algorithmic, process-based explanation, which should then be subject to rigorous empirical test. The abovementioned list of algorithmic level proposals inspired by Bayesian rational analysis, both in the domain of reasoning and in cognitive science more generally, gives grounds for reassurance. Moreover, the extensive empirical testing of these models (Green & Over 1997; 2000; McKenzie & Mikkelsen 2000; 2007; McKenzie et al. 2001; Nelson 2005; Oaksford & Moussakowski 2004; Oaksford & Wakefield 2003; Oaksford et al. 1999; 2000; Tenenbaum 1999) should allay concerns that rational analysis provides no testable predictions. Ironically, the only theories in the psychology of reasoning that have been algorithmically specified, aside from those within the Bayesian tradition, are directly based on another rational level theory: logic (Johnson-Laird 1992; Rips 1994). Theorists who have instead focused primarily on heuristics for reasoning have couched their explanations in purely verbal terms (Evans 1989; Evans & Over 2004). This indicates, we suggest, that rational analysis assists, rather than impedes, algorithmic explanation.

R4.4. Relevance of neural implementation

Bayesian rational analysis is, moreover, appealing because it appears to yield algorithms that can be implemented in the brain. In *BR* (Ch. 4), we observed that the Bayesian approach was sweeping across cognitive psychology. We might also have added that its influence in computational neuroscience is at least as significant (Friston 2005). Although our Bayesian analyses of higher-level reasoning do not directly imply Bayesian implementations at the algorithmic level, it is intriguing that influential theorists (Doya et al. 2007; Friston 2005; Ma et al. 2006) view Bayesian inference as providing the driving computational principle for neural information processing. Such models, using population codes (Ma et al. 2006), which avoid treating the brain as representing probabilities directly on a numerical scale, can model simple perceptual decision tasks (Gold & Shadlen 2000). Such convergence raises the possibility that Bayesian rational analyses of reasoning may one day find rather direct neural implementations.

De Neys specifically appeals to the implementation level in commenting on *BR*. He draws attention to imaging studies of reasoning that suggest a role for the anterior cingulate cortex in detecting conflict and inhibiting responses. As we have seen (sect. R3.4), such a role is entirely consistent with Bayesian approaches. Indeed, more broadly, imaging work on human reasoning, pioneered by Goel (e.g., Goel 2007), is at an exploratory stage, and currently provides few constraints on theory. Moreover, as we have seen, where cognitive neuroscientists concentrate on *what* computations the brain performs rather than *where*, the emerging answer is Bayesian.

R4.5. Optimality and rational analysis

A range of commentators (e.g., **Brighton & Olsson**, **Danks & Eberhardt**, **Evans**, and **Schroyens**) argue that the methodology of rational analysis faces conceptual problems. Our general response to these concerns is pragmatic. As with any methodology, we see rational analysis, using probabilistic methods or otherwise, as primarily to be judged by its results. Anderson's path-breaking work (1990; 1991a), and the huge literature on Bayesian models across the brain and cognitive sciences, of which *BR* is a part, is therefore, in our view, the best argument for the value of the approach. Parallels with closely related work in behavioral ecology and rational choice explanation in economics give further weight to the view that a "rational" style of explanation can yield considerable insights. But, like any style of explanation, rational analysis has its limits. Just as, in biology, some behaviors or structures are products of "history" rather than adaptation (Carroll 2005), and some economic behaviors are the product of cognitive limitations (e.g., Ariely et al. 2003; Thaler 2005), so in the brain and cognitive sciences, we should expect some phenomena to arise from specific aspects of algorithms/representations or neural implementation.

We are therefore happy to agree with commentators who suggest that there are cognitive phenomena for which purely rational considerations provide an incomplete, or indeed incorrect, explanation (e.g., **Brighton & Olsson**, **Evans**). We also agree that rational analysis is challenged where there are many, perhaps very different, near-optimal rational solutions (Brighton & Olsson). In such situations, rational analysis provides, at best, a range of options – but it does not provide an explanation of why one has been chosen. Nonetheless, these issues often cause few problems in practice, as the results in *BR* and in the wider program of rational explanation illustrate.

We agree, moreover, with concerns that finding *exactly* the optimal solution may be over-restrictive (**Brighton & Olsson**, **Evans**). Consider the case of perceptual organization, where the cognitive system must decide between multiple interpretations of a stimulus (Gregory 1970; von Helmholtz 1910/1925). Accounts based on Bayesian probability and on the closely related idea of maximizing simplicity (Chater 1996; Hochberg & McAlister 1953; Leeuwenberg & Boselie 1988) adopt the perspective of rational analysis, but they do so *comparatively*. That is, the perceptual system is presumed to choose interpretation A, rather than interpretation B, if A is more likely than B (or, in simplicity-based formulations, if it provides a *simpler* encoding of the sensory input). Neither the likelihood nor the simplicity principles in perceptual organization are presumed to imply that the perceptual system can optimize likelihood/simplicity – and indeed, in the general case, this is provably impossible (see Chater 1996, for discussion). Indeed, we suspect that rational analysis will, in many cases, primarily be concerned with providing a measure of the relative "goodness" of different cognitive processes or behaviors; and it is explanatory to the degree to which the "good" mechanisms are more prevalent than the "bad." The parallel with evolutionary explanation seems to be exact here: Inclusive fitness provides a crucial explanatory measure in explaining the evolution of biological structures, but the explanatory "bite" is comparative (i.e., in a certain environment, a flipper yields greater fitness than a leg). There is no assumption that biological evolution, in any context, reaches a state of completely optimized perfection; indeed, quite the reverse (Jacob 1977). Thus, Evans' emphasis on satisficing rather than optimizing, and Brighton & Olsson's focus on relative rationality, seem to us entirely consistent with *BR*.

Note, too, that in modeling many aspects of cognition, a full-scale rational analysis (specifying a task, environment, and computational limitations) may not be required. For example, conditional inference can be modeled in Bayesian terms, assuming only a probabilistic interpretation of the premises, and the requirement of maintaining consistent degrees of belief. The success of the probabilistic, rather than a logical, interpretation of the premises can be assessed by comparing the predictions of both approaches to data on human reasoning, as well general philosophical principles.

Brighton & Olsson also raise a different concern: that the specific sets of probabilistic assumptions (such as the independence assumptions embodied in naïve Bayes) may sometimes be justified *not* by rational analysis, but instead in the light of their general, formal properties, combined with empirical success in solving some externally defined task (e.g., estimating the relative sizes of German cities, Gigerenzer & Goldstein 1996). For example, a model such as naïve Bayes, they note, may be effective because it has few parameters and hence avoids over-fitting. We suggest, however, that this is not a separate type of explanation of inferential success, distinct from Bayesian rational analysis. Instead, the justification for preferring simple models can, itself, be provided in terms of Bayesian reasoning, and closely related formalisms, including minimum description length (Chater & Oaksford 2008b; MacKay 2003; Rissanen 1989; Vitányi & Li 2000).

R4.6. Need rational explanation be causal?

Brighton & Olsson, together with **Danks & Eberhardt**, raise the fundamental concern that rational explanation does not provide a *causal* explanation of behavior. We agree. Rational explanation is teleological (Fodor 1968) – it explains by reference to purpose, rather than cause.

In particular, rational explanation does not require that the rational analysis is itself represented in the mind of the agent, and does not, therefore, imply that behavior is governed by any such representation. Aerodynamics may provide an optimality-based explanation of the shape of the bird's wing; but aerodynamic calculations by the bird (or any other agent) are not causally responsible for the wing's shape.

Similarly, delineating the circumstances in which algorithms such as naïve Bayes (**Brighton & Olsson**; Domingos & Pazzani 1997), Take the Best (Gigerenzer & Goldstein 1996; Martignon & Hoffrage 1999), or unit-weighted regression (Dawes 1979) are reliable may require highly sophisticated rational explanation. Yet a cognitive system that employs such models may know nothing of such rational explanations – and indeed, these rational assumptions typically play no causal role in determining the behavior. Thus, in behavioral ecology, for example, the strategies animals use in foraging, mate selection, and so on, are typically explained using optimality explanations; but animals are not assumed to carry out optimality calculations to validate their behavioral strategies.

Danks & Eberhardt suggest that there is a "requirement for a teleological explanation that the normative principle must have played a causal role – ontogenetic, phylogenetic, or both – in the behavior's existence or persistence. 'Origin stories' are required for teleological explanation." We find this claim puzzling: normative principles, and rational explanations in general, are abstract – they are not part of the causal realm. Thus, a Bayesian rational analysis can no more *cause* a particular piece of behavior or reasoning, than the principles of arithmetic *cause* a calculator to display a particular number. Teleological explanations are distinctively non-causal, and necessarily so.

In this section, we have considered concerns about the general project of rational analysis. We now turn to consider specific issues relating to the rational models and empirical data presented in *BR*.

R5. Reconsidering models and data

Even if the broad sweep of arguments from the preceding sections is endorsed, there remain doubts about the details of the particular models described in *BR* and their ability to account for human reasoning data. Indeed, in the commentaries, issues of detail emerge most often between researchers who otherwise are in broad agreement. It is in this light that we consider the comments of **Liu**, **Oberauer**, **Over & Hadjichristidis**, and **Wagenmakers**. We also consider here **Halford**'s comments on syllogistic reasoning, drawn from a different framework.

R5.1. Conditional inference

Liu, **Oberauer**, and **Over & Hadjichristidis**, who have also advocated a probabilistic approach (in particular, to conditional inference), have concerns about our specific model. We addressed, in section R3.6, Over & Hadjichristidis's argument that we are not Bayesian enough, and that we should employ Jeffrey's rule to deal with uncertainty in the categorical premise of conditional inference. We pointed out that we too explicitly adopted Jeffrey's rule in *BR*. They also cite some unpublished results apparently showing that people have an imperfect understanding of Jeffrey's rule. These results are intriguing and suggest that more extensive empirical testing of this rule is required.[3]

Oberauer argues that our models of conditional inference and data selection may lead to absurdity. He argues that if the marginals, $P(p)$ and $P(q)$, remain fixed, which he describes as "axiomatic" in our theory,[4] then if one increases the probability that someone gets a headache, given they take drug X, then those who don't take X will get fewer headaches. This apparent absurdity stems from a conflation in Oberauer's description between the factual and the epistemic/doxastic: Changing this conditional *degree of belief* does not mean that these people *actually* achieve these benefits. In ignorance of the real conditional probability, but knowing the values of the marginals, I *should* revise my *degree of belief* that not taking this drug leads to fewer headaches. Yet this will only be appropriate when the marginals are known – which is clearly inappropriate in Oberauer's example.

Oberauer also perceives an inconsistency between our adoption of The Equation – $P(\textit{if } p \textit{ then } q) = P(q|p)$ – and our use of a contingency table to represent the conditional hypothesis in data selection. However, by The Equation there is only sufficient information in the premises of a conditional inference to draw MP by Bayesian (or Jeffrey) conditionalization (at least a point value). The remaining inferences can only be drawn on the assumption that people use the marginals to calculate the relevant conditional probabilities, for example, $P(\neg q|\neg p)$ for Denying the Antecedent (DA). Once $P(q|p)$ and the marginals are fixed, the contingency table is determined. Knowing the meaning of a statement is often equated with knowing the inferences that a statement licenses (Dowty et al. 1981). According to The Equation, the conditional only licenses "probabilized" MP. Probabilistically, to draw further inferences requires more information to be drawn from world knowledge. Hence, there is no inconsistency. Moreover, in the selection task, people are presented with an array of possible evidence types that makes the marginals relevant in the same way as presenting more than just MP in the conditional inference task. The degree of belief that is modified by selecting data is in the conditional *and the marginals*, which constitute the dependence and independence models. Thus, Oberauer's concerns can be readily addressed.

Oberauer also suggests that contingency tables are consistent with a probabilistic contrast approach, that is, the measure of the strength of an argument, for example, MP, is $P(q|p) - P(q|\neg p)$. It is for this reason that we believe that argument strength may indeed be two-dimensional (Oaksford & Hahn 2007). The conditional probability alone can mean that a good argument leads to no increase in the degree of belief in the conclusion, for example, for MP when $P(q|p) = P(q) = 1$. The probabilistic contrast (and other measures; see, e.g., **Nelson**, **Poletiek**, and Oaksford & Hahn 2007) captures the change in the probability of the conclusion brought about by an argument. Oberauer suggests that there is no evidence for people's use of the probabilistic contrast. Yet Over et al. (2007) found significant sensitivity to $P(q|\neg p)$, consistent with some use of the probabilistic contrast or a related measure of change, and the evidence is currently equivocal.

Oberauer also raises two concerns over evidence for our model of conditional inference. First, fitting a model with two free parameters to four data points "is no convincing accomplishment." Even so, as **Hahn** observes, the move to detailed model fitting of quantitative data represents significant progress in the psychology of reasoning (for early examples, see Krauth [1982] and Klauer [1999]). Moreover, in *BR* (pp. 146–49) we fitted the model to the 32 data points produced in Oaksford et al.'s (2000) Experiment 1 using only nine parameters, collapsing far more degrees of freedom than the model fitting reported in Oberauer (2006). Although Oberauer (2006) found poorer fits for our model than alternative theories, Oaksford and Chater (2008) found that the revised model presented in *BR* may provide better fits to Oberauer's data. Second, Oberauer argues that the most relevant empirical evidence comes from studies where probabilities were directly manipulated, of which he mentions two, Oaksford et al. (2000) and Oberauer et al. (2004). Moreover, he argues that their results are equivocal. However, several other studies have manipulated probabilities in conditional inference and found evidence in line with a probabilistic account (George 1997; Liu 2003; Liu et al. 1996; Stevenson & Over 1995). Oberauer also leaves aside the many studies on data selection showing probabilistic effects (see *BR*, Ch. 6).

Liu's arguments about second-order conditionalization point, we think, to an important factor that we have yet to consider in reasoning, that is, the effects of context. **Liu** has found that people often endorse the conclusion that, for example, *Tweety flies* on being told that *Tweety is a bird* in the absence of the conditional premise (reduced problems). This occurs because they fill in this information from world knowledge. However, **Liu** also found that endorsements increase when the conditional premise is added (complete problems). In *BR*, we argued that this occurs because people take the conditional premise as evidence that the conditional probability is higher (an inference that may arise from conversational pragmatics). **Liu** argues that our account implies that manipulations affecting reduced problems should also affect complete problems and provides evidence against this. Yet context, both cognitive and physical, may explain these differences in a way similar to recent studies of decision-making (Stewart et al. 2006). For example, suppose one is told about two swanneries, both containing the same number of swans. In one, 90% of swans are black (P(black|swan) = .9); in the other, 90% of swans are white (P(white|swan) = .9). On being told that *Tweety is a swan*, presumably one would only endorse *Tweety is white* at .5. This is because conversational pragmatics and world knowledge indicate that Tweety is in one of the just mentioned swanneries, but the dialogue up to this point does not indicate which one.[5] However, the addition of the conditional premise *if a bird is a swan it is white* immediately indicates which swannery is being talked about, that is, the one in which P(white|swan) is high, and now endorsements should increase to .9. Clearly, although manipulations of the relative number of swans in each swannery might affect the reduced problem, they should not affect the complete problem. So if the swannery in which most swans are black were one tenth of the size of the other swannery, then, given natural sampling assumptions, endorsements for the reduced problem should increase to .83, but endorsements of the complete problem should remain the same.

R5.2. Data selection

Wagenmakers raises a variety of concerns about our optimal data selection model. First, why do we concede that people should select the standard "logical" A card and 7 card choices, *if* the rule only applies to the four cards? In *BR* (p. 210), we argue that people rarely use conditionals to describe just four objects – they assume that the cards are drawn from a larger population. Consequently, we quite explicitly do not make the counterintuitive prediction that Wagenmakers ascribes to us. Second, Wagenmakers wonders why – when all cards carry some information – do participants not select all the cards, if they are maximizing information gain? We assume that the pragmatics of the task suggests to participants that they should select some cards, but not others (*BR*, pp. 200–201). Third, Wagenmakers suggests that incentivized individuals with more time might make the logical response. Work on individual differences (e.g., Stanovich & West 2000) is consistent with the view that logical competence is learned, either directly (e.g., studying logic or math) or indirectly (e.g., learning to program or learning conventional, non-Bayesian statistics); such logical competence is a prerequisite for "logical" responses, and covaries with IQ as measured in University populations. Wagenmakers also remarks that, as Bayesians, we should avoid null hypothesis testing in statistically assessing our models. This choice is purely pragmatic: it conforms to the current demands of most journals.

R5.3. Syllogisms and development

Halford argues that mental models theory and a relational complexity measure fit the data as well as the probability heuristics model (PHM), conceding, however, that only PHM generalizes to *most* and *few*. Copeland (2006) has also recently shown that PHM provides better fits than mental models and mental logic for extended syllogisms involving three quantified premises. Halford also suggests that basing confidence in the conclusion on the least probable premise, as in our *max*-heuristic, is counterintuitive. He proposes that confidence should instead be based on relational complexity, which covaries with the least probable premise. But perhaps Halford's intuition goes the wrong way: the least probable premise is the most informative; and surely the more information you are given, the stronger the conclusions you can draw?

De Neys and **Straubinger, Cokely, & Stevens** (**Straubinger et al.**) both argue that there are important classes of evidence that we do not address. De Neys argues that attention to latency data and imaging studies provides a greater role for logic, a claim we disputed earlier. Note, also, that the algorithmic theory in PHM has been applied to latency data and accounts for the data, as well as mental models (Copeland & Radvansky 2004). Straubinger et al. are concerned that we ignore developmental data. In particular, they view the findings on the development of working memory as providing a particular challenge to a Bayesian approach. They do, however, acknowledge that in different areas (e.g., causal reasoning), Bayesian ideas are being successfully applied to developmental data (Navarro et al. 2006; Sobel et al. 2004). Straubinger et al.'s emphasis on working memory provides good reason to believe that our particular approach to deductive reasoning may extend to development. Copeland and Radvansky (2004) explicitly related working-memory limitations to PHM, finding that it provided as good an explanation as mental models theory of the relationship between working-memory capacity and reasoning performance. This result provides some indication that, at least for syllogistic reasoning, developmental trajectories explicable by mental models may be similarly amenable to explanation in terms of probability heuristics. Our approach also provides a natural way in which experience, leading to the learning of environmental statistics, might influence reasoning development. Exploring these possibilities must await future research.

R6. The Bayesian turn

BR is part of a larger movement across the brain and cognitive sciences – a movement which sees cognition as centrally concerned with uncertainty; and views Bayesian probability as the appropriate machinery for dealing with uncertainty. Probabilistic ideas have become central to theories of elementary neural function (Doya et al. 2007), motor control (Körding & Wolpert 2004), perception (Knill & Richards 1996), language processing (Manning & Schütze 1999), and high-level cognition (Chater & Oaksford 2008a; Chater et al. 2006). They also cut across Marr's (1982) computational (Anderson 1990; Pearl 2000), algorithmic (Jacobs et al. 1991), and implementational (Doya et al. 2007) levels of explanation. In arguing that commonsense reasoning should be understood in terms of probability, we are merely recasting Laplace's (1814/1951) classic dictum concerning the nature of probability theory: "The theory of probabilities is at bottom nothing but common sense reduced to calculus."

NOTES

1. Although Braine and O'Brien (1991) explicitly reject the use of relevance logic (Anderson & Belnap 1975), this does provide an interesting possible route for developing these ideas. In particular, interpretations of the semantics of relevance logics as a ternary relation between possible worlds, or from an information-theoretic perspective, as a ternary relation between a source, a receiver, and a channel (Restall 1996), may provide interesting connections with nonmonotonic reasoning.

2. By contrast, we know of just one paper in the psychology of reasoning discussing Dempster-Shafer belief functions, namely, George (1997).

3. Its normative status has also been questioned for many years (see, e.g., Field 1978).

4. This is despite the fact that they were not fixed in Oaksford and Chater (1994).

5. Of course, different assumptions would yield different results. For example, if the previous dialogue had been talking about the swannery, where most swans are black, just before introducing Tweety, the assumption may be that Tweety comes from that swannery and so *Tweety is white* might only be endorsed at .1.

References

[The letters "a" and "r" before author's initials stand for target article and response references, respectively]

Adams, E. W. (1975) *The logic of conditionals: An application of probability to deductive logic*. Reidel. [arMO, NP]

(1998) *A primer of probability logic*. CLSI Publications. [arMO]

Ali, N., Schlottmann, A., Shaw, A., Chater, N. & Oaksford, M. (in press) Causal discounting and conditional reasoning in children. In: *Cognition and conditionals: Probability and logic in human thought*, ed. M. Oaksford & N. Chater. Oxford University Press. [rMO]

Anderson, A. & Belnap, N. D. (1975) *Entailment: The logic of relevance and necessity, vol. 1*. Princeton University Press. [rMO]

Anderson, J. R. (1983) *The architecture of cognition*. Harvard University Press. [rMO]

(1990) *The adaptive character of thought*. Erlbaum. [TLG, UH, ELK, CRMM, arMO, WS]

(1991a) Is human cognition adaptive? *Behavioral and Brain Sciences* 14:471–84; discussion 485–517. [HB, arMO, NS]

(1991b) The adaptive nature of human categorization. *Psychological Review* 98:409–29. [aMO]

Anderson, J. R. & Matessa, M. (1998) The rational analysis of categorization and the ACT-R architecture. In: *Rational models of cognition*, ed. M. Oaksford & N. Chater, pp.197–217. Oxford University Press. [aMO]

Anderson, J. R. & Milson, R. (1989) Human memory: An adaptive perspective. *Psychological Review* 96:703–19. [aMO]

Anderson, J. R. & Schooler, L. J. (1991) Reflections of the environment in memory. *Psychological Science* 2:396–408. [aMO]

Antoniou, G. (1997) *Nonmonotonic reasoning*. MIT Press. [KS]

Ariely, D., Loewenstein, G. & Prelec, D. (2003) Coherent arbitrariness: Stable demand curves without stable preferences. *Quarterly Journal of Economics* 118:73–105. [rMO]

Aristotle (1980) *Nicomachean ethics*, trans. W. D. Ross. Clarendon Press. [aMO]

Ball, L. J., Philips, P., Wade, C. N. & Quayle, J. D. (2006) Effects of belief and logic on syllogistic reasoning: Eye-movement evidence for selective processing models. *Experimental Psychology* 53:77–86. [WDN]

Baltes, P. B., Staudinger, U. M. & Lindenberger, U. (1999) Lifespan psychology: Theory and application to intellectual functioning. *Annual Review of Psychology* 50:471–507. [NS]

Baron, J. (1981) An analysis of confirmation bias. Paper presented at the 22nd Annual Meeting of the Psychonomic Society. 6–8 November, 1981, Philadelphia, PA. [aMO]

(1985) *Rationality and intelligence*. Cambridge University Press. [JDN, aMO]

(2004) Normative models of judgment and decision making. In: *Blackwell handbook of judgment and decision making*, ed. D. J. Koehler & N. Harvey, pp. 19–36. Blackwell. [JDN]

Barwise, J. & Cooper, R. (1981) Generalized quantifiers and natural language. *Linguistics and Philosophy* 4:159–219. [aMO]

Becker, G .S. (1976) *The economic approach to human behaviour*. University of Chicago Press. [ELK]

Benferhat, S., Bonnefon, J. F. & Da Silva Neves, R. (2005) An overview of possibilistic handling of default reasoning, with experimental studies. *Synthese* 146:53–70. [GP]

Bennett, J. (2003) *A philosophical guide to conditionals*. Oxford University Press. [arMO]

Bernoulli, J. (1713/2005) *Ars conjectandi* [*The art of conjecture*], trans. E. D. Sylla. Johns Hopkins University Press. [aMO]

Biazzo, V., Gilio, A., Lukasiewicz, T. & Sanfilippo, G. (2005) Probabilistic logic under coherence: Complexity and algorithms. *Annals of Mathematics and Artificial Intelligence* 45(1–2):35–81. [NP]

Blakemore, C., Adler, K. & Pointon, M. (1990) *Vision: Coding and efficiency*. Cambridge University Press. [rMO]

Bookstaber, R. & Langsam, J. (1985) On the optimality of coarse behavior rules. *Journal of Theoretical Biology* 116:161–93. [HB]

Boole, G. (1854/1958) *An investigation of the laws of though*t. Macmillan/Dover. (Reprinted by Dover, 1958). [aMO]

Bovens, L. & Hartmann, S. (2003) *Bayesian epistemology*. Clarendon Press. [aMO]

Braine, M. D. S. (1978) On the relation between the natural logic of reasoning and standard logic. *Psychological Review* 85:1–21. [aMO]

(1990) The natural logic approach to reasoning. In: *Reasoning, necessity, and logic: Developmental perspectives*, ed. W. F. Overton. Erlbaum. [DPO]

Braine, M. D. S. & O'Brien, D. P. (1991) A theory of *if*: A lexical entry, reasoning program, and pragmatic principles. *Psychological Review* 98:182–203. [DPO, rMO]

Braine, M. D. S. & O'Brien, D. P., eds. (1998) *Mental logic*. Erlbaum. [DPO]

Brighton, H. & Gigerenzer, G. (2008) Bayesian brains and cognitive mechanisms: Harmony or dissonance? In: *The probabilistic mind: Prospects for Bayesian cognitive science*, ed. N. Chater & M. Oaksford, pp. 189–208, Oxford University Press. [HB, NS]

Bullinaria, J. A. & Chater, N. (1995) Connectionist modelling: Implications for neuropsychology. *Language and Cognitive Processes* 10:227–64. [UH]

Byrne, R. M. J. (1989) Suppressing valid inferences with conditionals. *Cognition* 31:61–83. [KS]

Byrne, R. M. J., Espino, O. & Santamaria, C. (1999) Counterexamples and the suppression of inferences. *Journal of Memory and Language* 40:347–73. [I-mL]

Carnap, R. (1950) *The logical foundations of probability*. University of Chicago Press. [rMO, FHP]

Carroll, S. (2005) *Endless forms most beautiful*. W. W. Norton. [rMO]

Chater, N. (1995) Neural networks: The new statistical models of mind. In: *Connectionist models of memory and language*, ed. J. P. Levy, D. Bairaktaris,

J. A. Bullinaria & P. Cairns, pp. 207–27. University College London Press. [UH]
(1996) Reconciling simplicity and likelihood principles in perceptual organization. *Psychological Review* 103:566–81. [rMO]
Chater, N., Crocker, M. & Pickering, M. (1998) The rational analysis of inquiry: The case of parsing: In: *Rational models of cognition*, ed. M. Oaksford & N. Chater, pp. 441–68. Oxford University Press. [aMO]
Chater, N. & Manning, C. D. (2006) Probabilistic models of language processing and acquisition. *Trends in Cognitive Sciences* 10:335–44. [TLG, aMO]
Chater, N. & Oaksford, M. (1999a) Ten years of the rational analysis of cognition. *Trends in Cognitive Science* 3:57–65. [TLG]
(1999b) The probability heuristics model of syllogistic reasoning. *Cognitive Psychology* 38:191–258. [arMO]
(2006) Mental mechanisms: Speculations on human causal learning and reasoning. In: *Information sampling and adaptive cognition*, ed. K. Fiedler & P. Juslin, pp. 210–38. Cambridge University Press. [rMO]
eds. (2008a) *The probabilistic mind*: *Prospects for Bayesian cognitive science*. Oxford University Press. [UH, rMO]
(2008b) The probabilistic mind: Where next? In: *The probabilistic mind: Prospects for Bayesian cognitive science*, ed. N. Chater & M. Oaksford, pp. 501–14. Oxford University Press. [rMO]
Chater, N., Oaksford, M., Heit, E. & Hahn, U. (in press) Inductive logic and empirical psychology. In: *The handbook of philosophical logic, vol. 10*, ed. S. Hartmann & J. Woods. Springer. [rMO]
Chater, N., Oaksford, M., Nakisa, R. & Redington, M. (2003) Fast, frugal and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes* 90:63–86. [JDN, rMO, WS]
Chater, N., Tenenbaum, J. B. & Yuille, A. eds. (2006) Probabilistic models of cognition: Where next? *Trends in Cognitive Sciences* 10:335–44. [Special Issue.] [UH, rMO]
Chater, N. & Vitányi, P. (2002) Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences* 7:19–22. [rMO]
Cheng, P. W. & Holyoak, K. J. (1985) Pragmatic reasoning schemas. *Cognitive Psychology* 17:391–416. [aMO]
Chomsky, N. (1957) *Syntactic structures*. Mouton. [aMO]
(1965) *Aspects of the theory of syntax.* MIT Press. [aMO]
Chou, T. H. (2007) The mechanism of suppression effects in conditional reasoning. Unpublished Doctoral dissertation, National Chung-Cheng University, Department of Psychology, Chia-Yi, Taiwan. [I-mL]
Cohen, L. J. (1981) Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences* 4:317–70. [JDN, aMO]
Cokely, E. T., Kelley, C. M. & Gilchrist, A. L. (2006) Sources of individual differences in working memory: Contributions of strategy to capacity. *Psychonomic Bulletin and Review* 13:991–97. [NS]
Coletti, G. & Scozzafava, R. (2002) Probabilistic logic in a coherent setting. Kluwer. [NP]
Conlin, J. A., Gathercole, S. E. & Adams, J. W. (2005) Children's working memory: Investigating performance limitations in complex span tasks. *Journal of Experimental Child Psychology* 90:303–17. [NS]
Cooper, G. F. (1990) The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence* 42:393–405. [HB]
Copeland, D. E. (2006) Theories of categorical reasoning and extended syllogisms. *Thinking and Reasoning* 12:379–412. [rMO]
Copeland, D. E. & Radvansky, G. A. (2004) Working memory and syllogistic reasoning. *Quarterly Journal of Experimental Psychology* 57A:1437–57. [arMO]
Cosmides, L. (1989) The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition* 31:187–276. [arMO]
Cosmides, L. & Tooby, J. (1996) Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition* 58:1–73. [JDN]
(2000) Evolutionary psychology and the emotions. In: *Handbook of emotions*, 2nd edition, ed. M. Lewis & J. M. Haviland-Jones, pp. 91–115. Guilford. [aMO]
Courville, A. C., Daw, N. D. & Touretzky, D. S. (2006) Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences* 10:294–300. [aMO]
Damer, E. T. (2005) *Attacking faulty reasoning: A practical guide to fallacy-free arguments, 5th edition*. Thomson Wadsworth. [WS]
Danks, D. (2008) Rational analyses, instrumentalism, and implementations. In: *The probabilistic mind: Prospects for Bayesian cognitive science*, ed. N. Chater & M. Oaksford. Oxford University Press. [HB]
Daston, L. (1988) *Classical probability in the enlightenment*. Princeton University Press. [aMO]
Davidson, D. (1984) *Inquiries into truth and interpretation*. Oxford University Press. [aMO]
Dawes, R. M. (1979) The robust beauty of improper linear models in decision making. *American Psychologist* 34:571–82. [rMO]
de Finetti, B. (1974/1975) *Theory of probability, vols. 1 & 2*. Wiley. [NP]
Dempster, A. P. (1968) A generalization of Bayesian inference. *Journal of the Royal Statistical Society, Series B* 30:205–47. [rMO]
Dennett, D. (1987) *The intentional stance*. MIT Press. [WS]
De Neys, W. & Franssens, S. (2007) The nature of belief inhibition during thinking: How reasoning impairs memory. *Proceedings of the Annual Meeting of the Cognitive Science Society* 29:8. [WDN]
De Neys, W. & Glumicic, T. (2008) Conflict monitoring in dual process theories of thinking. *Cognition* 106:1248–99. [WDN]
De Neys, W., Vartanian, O. & Goel, V. (2008) Smarter than we think: When our brains detect that we are biased. *Psychological Science* 19:483–89. [WDN, rMO]
Dick, F., Leech, R., Moses, P. & Saccuman, M. C. (2006) The interplay of learning and development in shaping neural organization. *Developmental Science* 9:14–17. [NS]
Domingos, P. & Pazzani, M. (1997) On the optimality of the simple Bayesian classifier under zero-one loss. *Machine Learning* 29:103–30. [HB, rMO]
Dowty, D. R., Wall, R. E. & Peters, S. (1981) *Introduction to Montague semantics*. Springer. [rMO]
Doya, K., Ishii, S., Rao, R. P. N. & Pouget, A., eds. (2007) *The Bayesian brain: Probabilistic approaches to neural coding.* MIT Press. [rMO]
Dubois, D. & Prade, H. (1988) *Possibility theory.* Plenum. [GP]
Earman, J. (1992) *Bayes or bust*? MIT Press. [arMO]
Edgington, D. (1995) On conditionals. *Mind* 104:235–329. [UH, arMO]
Elqayam, S. (2008) Normative rationality and the is-ought fallacy. In: *Proceedings of the Second European Cognitive Science Conference*, ed. S. Vosiniadou, D. Kayser & A. Protopaps, pp. 249–99. Psychology Press. [JStBTE]
Ericsson, K. A., Charness, N., Hoffmann, R. R. & Feltovich, P. J. (2006) *Cambridge handbook of expertise and expert performance.* Cambridge University Press. [NS]
Ericsson, K. A. & Kintsch, W. (1995) Long-term working memory. *Psychological Review* 102:211–45. [NS]
Ericsson, K. A. & Lehmann, A. C. (1996) Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Reviews in Psychology* 47:273–305. [NS]
Evans, J. St. B. T. (1972) Reasoning with negatives. *British Journal of Psychology* 63:213–19. [aMO]
(1977) Linguistic factors in reasoning. *Quarterly Journal of Experimental Psychology* 29A:297–306. [I-mL]
(1989) *Bias in human reasoning: Causes and consequences*. Erlbaum. [rMO]
(2002) Logic and human reasoning: An assessment of the deduction paradigm. *Psychological Bulletin* 128:978–96. [JStBTE]
(2007) *Hypothetical thinking: Dual processes in reasoning and judgement*. Psychology Press. [JStBTE, rMO]
(2008) Dual-processing accounts of reasoning, judgment and social cognition. *Annual Review of Psychology* 59:255–78. [JStBTE]
Evans, J. St. B. T. & Frankish, K., eds. (in press) *In two minds: Dual processes and beyond*. Oxford University Press. [rMO]
Evans, J. St. B. T. & Handley, S. J. (1999) The role of negation in conditional inference. *Quarterly Journal of Experimental Psychology* 52A:739–69. [KO, aMO]
Evans, J. St. B. T., Handley, S. J., Harper, C. N. J. & Johnson-Laird, P. N. (1999) Reasoning about necessity and possibility: A test of the mental model theory of deduction. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25:1495–1513. [aMO]
Evans, J. St .B. T., Handley, S. J. & Over, D. E. (2003) Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory and Cognition* 29321–55. [I-mL, aMO]
Evans, J. St. B. T. & Lynch, J. S. (1973) Matching bias in the selection task. *British Journal of Psychology* 64:391–97. [aMO]
Evans, J. St. B. T., Newstead, S. E. & Byrne, R. J. (1993) *Human reasoning.* Erlbaum. [aMO]
Evans, J. St. B. T. & Over, D. E. (1996a) *Rationality and reasoning*. Psychology Press. [JStBTE, DEO, arMO]
(1996b) Rationality in the selection task: Epistemic utility versus uncertainty reduction. *Psychological Review* 103:356–63. [aMO]
(2004) *If*. Oxford University Press. [JStBTE, UH, arMO, NP]
Evans, J. St. B. T., Stevenson, R. J., Over, D. E., Handley, S. J. & Sloman, S. A. (2004) *ESRC end of grant report*, No. R00239074. The Economic and Social Research Council. [DEO]
Falk, R. & Wilkening, F. (1998) Children's construction of fair chances: Adjusting probabilities. *Developmental Psychology* 34:1340–57. [NS]
Field, H. (1978) A note on Jeffrey conditionalization. *Philosophy of Science* 45:361–67. [rMO]
Fodor, J. A. (1968) *Psychological explanation*. Random House. [rMO]
(1983) *The modularity of mind*. MIT Press. [UH, aMO]

Fodor, J. A. & Pylyshyn, Z. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition* 28:183–204. [UH]

Ford, M. (2004) System LS: A three tiered nonmonotonic reasoning system. *Computational Intelligence* 20:89–108. [GP]

Franks, B. (1995) On explanation in the cognitive sciences: Competence, idealization, and the failure of the classical cascade. *British Journal of the Philosophy of Science* 46(4):475–502. [WS]

(1999) Idealizations, competence and explanation: A response to Patterson. *British Journal of the Philosophy of Science* 50:735–46. [WS]

Friedman, M. (1953) *Essays in positive economics.* University of Chicago Press. [ELK]

Friedman, N., Getoor, L., Koller, D. & Pfeffer, A. (1999) Learning probabilistic relational models. In: *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI)*, ed. T. Dean, pp. 1300–309. Morgan Kaufmann. [TLG]

Friedman, N. & Halpern, J. Y. (1995) Plausibility measures: A user's guide. In: *Proceedings of the eleventh conference on uncertainty in AI*, ed. P. Besnard & S. Hanks, pp. 175–84. Morgan Kaufmann. [GP]

Friston, K. (2005) A theory of cortical responses. *Philosophical Transactions of the Royal Society B* 360:815–36. [rMO]

Galotti, K. M., Komatsu, L. K. & Voelz, S. (1997) Children's differential performance on deductive and inductive syllogisms. *Developmental Psychology* 33:70–78. [NS]

Gärdenfors, P. (1986) Belief revisions and the Ramsey test for conditionals. *Philosophical Review* 95:81–93. [rMO]

George, C. (1997) Reasoning from uncertain premises. *Thinking and Reasoning* 3:161–90. [rMO]

Geurts, B. (2003) Reasoning with quantifiers. *Cognition* 86:223–51. [aMO]

Gigerenzer, G. (1996) On narrow norms and vague heuristics: A reply to Kahneman and Tversky. *Psychological Review* 103:592–96. [CRMM]

(2004) Fast and frugal heuristics: The tools of bounded rationality. In: *Blackwell handbook of judgment and decision making*, ed. D. J. Koehler & N. Harvey, pp. 62–88. Blackwell. [JStBTE]

Gigerenzer, G. & Goldstein, D. (1996) Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103:650–69. [arMO]

Gigerenzer, G. & Hoffrage, U. (1995) How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review* 102(4):684–704. [JDN, aMO]

Gigerenzer, G., Swijinck, Z., Porter, T., Daston, L., Beatty, J. & Kruger, L. (1989) *The empire of chance*. Cambridge University Press. [aMO]

Gigerenzer, G., Todd, P. & the ABC Research Group. (1999) *Simple heuristics that make us smart*. Oxford University Press. [HB, arMO]

Gilio, A. (2002) Probabilistic reasoning under coherence in System P. *Annals of Mathematics and Artificial Intelligence* 34:5–34. [NP]

Gilovich, T., Griffin, D. & Kahneman, D. (2002) *Heuristics and biases: The psychology of intuitive judgement*. Cambridge University Press. [JStBTE, CRMM]

Glymour, C. (2001) *The mind's arrow*. MIT Press. [rMO]

Goel, V. (2007) The anatomy of deduction. *Trends in Cognitive Science* 11:435–41. [rMO]

Gold, J. I. & Shadlen, M. N. (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–94. [rMO]

Good, I. J. (1950) *Probability and the weighing of evidence*. Griffin. [JDN]

(1975) Explicativity, corroboration, and the relative odds of hypotheses. *Synthese* 30:39–73. [JDN]

Green. D. W. & Over, D. E. (1997) Causal inference, contingency tables and the selection task. *Current Psychology of Cognition* 16:459–87. [arMO]

(2000) Decision theoretical effects in testing a causal conditional. *Current Psychology of Cognition* 19:51–68. [arMO]

Gregory, R. L. (1970) *The intelligent eye*. Weidenfeld & Nicolson. [rMO]

Grice, H. P. (1975) Logic and conversation. In: *The logic of grammar*, ed. D. Davidson & G. Harman, pp. 64–75. Dickenson. [rMO]

Griffin, D. R. (1992) *Animal minds*. University of Chicago Press. [ELK]

Griffiths, T. L. & Ghahramani, Z. (2006) Infinite latent feature models and the Indian buffet process. In: *Advances in neural information processing systems, vol. 18*, ed. Y. Weiss, B. Scholkopf & J. Plaut, pp. 475–82. MIT Press. [TLG]

Griffiths, T. L., Steyvers, M. & Tenenbaum, J. B. (2007) Topics in semantic representation. *Psychological Review* 114:211–44. [rMO]

Griffiths, T. L. & Tenenbaum, J. B. (2005) Structure and strength in causal induction. *Cognitive Psychology* 51:354–84. [TLG, arMO]

Hacking, I. (1975) *The emergence of probability*. Cambridge University Press. [aMO]

(1990) *The taming of chance*. Cambridge University Press. [aMO]

Hadjichristidis, C., Sloman, S. A. & Over, D. E. (in preparation) Revision of beliefs from probabilistic categorical arguments. [DEO]

Hahn, U. & Oaksford, M. (2007) The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review* 114:704–32. [rMO]

Halford, G. S., Cowan, N. & Andrews, G. (2007) Separating cognitive capacity from knowledge: A new hypothesis. *Trends in Cognitive Sciences* 11(6):236–42. [GSH]

Halpern, J. Y. (1997) Defining relative likelihood in partially-ordered preferential structures. *Journal of Artificial Intelligence Research* 7:1–24. [GP]

Harman, G. (1986) *Change in view: Principles of reasoning*. MIT Press. [rMO]

Hattori, M. (2002) A quantitative model of optimal data selection in Wason's selection task. *Quarterly Journal of Experimental Psychology* 55A:1241–72. [aMO]

Hawks, J., Wang, E. T., Cochran, G. M., Harpending, H. C. & Moyzis, R. K. (2007) Recent acceleration of human adaptive evolution. *Proceedings of the National Academy of Sciences USA* 104(52):20753–58. [JDN]

Hawthorn, J. (2008) Inductive logic. In: *Stanford Encyclopedia of Philosophy*. Available at: http://plato.stanford.edu/entries/logic-inductive/. [rMO]

Hawthorne, J. & Makinson, D. (2007) The quantitative/qualitative watershed for rules of uncertain inference. *Studia Logica* 86:247–97. [NP]

Henle, M. (1978) Foreword. In: *Human reasoning*, ed. R. Revlin & R. E. Mayer. Winston. [aMO]

Hilton, D. J., Kemmelmeier, M. & Bonnefon, J-F. (2005) Putting Ifs to work: Goal-based relevance in conditional directives. *Journal of Experimental Psychology: General* 134:388–405. [rMO]

Hochberg, J. & McAlister, E. (1953) A quantitative approach to figure "goodness." *Journal of Experimental Psychology* 46:361–64. [rMO]

Hoffman, D. (in press) The interface theory of perception: Natural selection drives true perception to swift extinction. In: *Object categorization: Computer and human vision perspectives*, ed. S. Dickinson, M. Tarr, A. Leonardis & B. Schiele. Cambridge University Press. [JDN]

Hoffrage, U. (2000) Why the analyses of cognitive processes matter. *Behavioral and Brain Sciences* 23:679–80. [WDN]

Horwich, P. (1982) *Probability and evidence*. Cambridge University Press. [aMO]

Houdé, O., Zago, L., Mellet, E., Moutier, S., Pineau, A., Mazoyer, B. & Tzourio-Mazoyer, N. (2000) Shifting from the perceptual brain to the logical brain: The neural impact of cognitive inhibition training. *Journal of Cognitive Neuroscience* 12:721–28. [WDN, rMO]

Howe, M. L. & Rabinowitz, F. M. (1996) Reasoning from memory: A lifespan inquiry into the necessity of remembering when reasoning about class inclusion. *Journal of Experimental Child Psychology* 61:1–42. [NS]

Howson, C. & Urbach, P. (1993) *Scientific reasoning: The Bayesian approach*, 2nd edition. Open Court. [DEO, aMO]

Hurley, S. & Nudds, M., eds. (2006) *Rational animals?* Oxford University Press. [rMO, ELK]

Inhelder, B. & Piaget, J. (1955) *De la logique de l'enfant à la logique de l'adolescent.* Presses Universitaires de France. (English version: *The growth of logical thinking from childhood to adolescence.* Routledge, 1958). [aMO]

Jacob F. (1977) Evolution and tinkering. *Science* 196:1161–66. [rMO]

Jacobs, R. A., Jordan, M. I., Nowlan, S. & Hinton, G. E. (1991) Adaptive mixtures of local experts. *Neural Computation* 3:1–12. [rMO]

Jeffrey, R. C. (1967) *Formal logic: Its scope and limits*, 2nd edition. McGraw-Hill. [rMO]

(1983) *The logic of decision*, 2nd edition. University of Chicago Press. [DEO, arMO]

Johnson-Laird, P. N. (1983) *Mental models*. Cambridge University Press. [arMO]

(1992) *Syllogs* (computer program). Available at: http://webscript.princeton.edu/~mentmod/models.php. [rMO]

Johnson-Laird, P. N. & Byrne, R. M. J. (1991) *Deduction.* Erlbaum. [aMO]

(2002) Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review* 109:646–78. [aMO]

Juhl, C. F. (1993) Bayesianism and reliable scientific inquiry. *Philosophy of Science* 60:302–19. [DD]

Kahneman, D., Slovic, P. & Tversky, A., eds. (1982) *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press. [CRMM, aMO]

Kahneman, D. & Tversky, A. (2000) *Choices, values, and frames*. Cambridge University Press. [CRMM]

Kakade, S. & Dayan, P. (2002) Acquisition and extinction in autoshaping. *Psychological Review* 109:533–44. [aMO]

Kant, E. (1787/1961) *Critique of pure reason*, *First edition, second impression*, trans. N. K. Smith. Macmillan. [aMO]

Kearns, M., Mansour, Y., Ng, A. Y. & Ron, D. (1997) An experimental and theoretical comparison of model selection methods. *Machine Learning* 27:7–50. [HB]

Kemp, C. & Tenenbaum, J. B. (2008) The discovery of structural form. *Proceedings of the National Academy of Sciences USA* 105:10687–92. [rMO]

Kirby, K. N. (1994) Probabilities and utilities of fictional outcomes in Wason's four card selection task. *Cognition* 51:1–28. [aMO]

Klauer, K. C. (1999) On the normative justification for information gain in Wason's selection task. *Psychological Review* 106:215–22. [arMO]

Klauer, K. C., Stahl, C. & Erdfelder, E. (2007) The abstract selection task: New data and an almost comprehensive model. *Journal of Experimental Psychology: Learning, Memory and Cognition* 33:680–703. [UH]
Klayman, J. & Ha, Y.-W. (1987) Confirmation, disconfirmation, and information. *Psychological Review* 94:211–28. [JDN]
Kneale, W. & Kneale, M. (1962) *The development of logic*. Oxford University Press. [DPO]
Knill, D. & Richards, W., eds. (1996) *Perception as Bayesian inference.* Cambridge University Press. [arMO]
Koehler, D. J. & Harvey, N. (2004) *Blackwell handbook of judgement and decision making*. Blackwell. [JStBTE]
Körding, K. P. & Wolpert, D. (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–47. [rMO]
Kowalski, R. (1979) Algorithm = Logic + Control. *Communications of the Association for Computing Machinery* 22:424–36. [rMO]
Kraus, S., Lehmann, D. & Magidor, M. (1990) Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167–207. [NP]
Krauth, J. (1982) Formulation and experimental verification of models in propositional reasoning. *Quarterly Journal of Experimental Psychology* 34:285–98. [rMO]
Kreps, D. M. (1990) *A course in microeconomic theory*. Princeton University Press. [ELK]
Krueger, J. I. & Funder, D. C. (2004) Towards a balanced social psychology: Causes, consequences, and cures for the problem-seeking approach to social behavior and cognition. *Behavioral and Brain Sciences* 27(3):313–76. [CRMM]
Kuhn, T. (1962) *The structure of scientific revolutions*. University of Chicago Press. [aMO]
Kuncheva, L. I. (2006) On the optimality of Naïve Bayes with dependent binary features. *Pattern Recognition Letters* 27:830–37. [HB]
Kyllonen, P. C. & Christal, R. E. (1990) Reasoning ability is (little more than) working-memory capacity?! *Intelligence* 14:389–433. [NS]
Lakatos, I. (1970) Falsification and the methodology of scientific research programmes. In: *Criticism and the growth of knowledge*, ed. I. Lakatos & A. Musgrave, pp. 91–196. Cambridge University Press. [aMO]
Laplace, P. S. (1951) *A philosophical essay on probabilities*, trans. F. W. Truscott & F. L. Emory. Dover. (Original work published 1814). [rMO]
Lauritzen, S. & Spiegelhalter, D. (1988) Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society B* 50:157–224. [rMO]
Leeuwenberg, E. & Boselie, F. (1988) Against the likelihood principle in visual form perception. *Psychological Review* 95:485–91. [rMO]
Levi, I. (1988) The demons of decision. *Monist* 70:193–211. [DD]
(2002) Money pumps and diachronic books. *Philosophy of Science* 69:S235–47. [DD]
Lindley, D. V. (1956) On a measure of the information provided by an experiment. *Annals of Mathematical Statistics* 27:986–1005. [JDN, aMO]
(2006) *Understanding uncertainty*. Wiley. [NP]
Lipsey, R. G. & Lancaster, K. (1956) The general theory of second best. *Review of Economic Studies* 24:11–32. [HB]
Liu, I.-M. (2003) Conditional reasoning and conditionalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29:694–709. [I-mL, rMO]
Liu, I.-M. & Chou, T. H. (2008) A developmental analysis of conditional reasoning. Unpublished manuscript, National Chung-Cheng University, Chia-Yi, Taiwan. [I-mL]
Liu, I. M., Lo, K. C. & Wu, J. T. (1996) A probabilistic interpretation of "If-Then". *The Quarterly Journal of Experimental Psychology* 49A:828–44. [I-mL, rMO]
Lukasiewicz, T. (2005) Weak nonmonotonic probabilistic logics. *Artificial Intelligence* 168:119–61. [NP]
Ma, W. J., Beck, J., Latham, P. & Pouget, A. (2006) Bayesian inference with probabilistic population codes. *Nature Neuroscience* 9:1432–38. [rMO]
MacKay, D. J. C. (2003) *Information theory, inference, and learning algorithms.* Cambridge University Press. [rMO]
Maher, P. (1992) Diachronic rationality. *Philosophy of Science* 59:120–41. [DD]
Manktelow, K. I. & Over, D. E. (1987) Reasoning and rationality. *Mind and Language* 2:199–219. [aMO]
(1991) Social roles and utilities in reasoning with deontic conditionals. *Cognition* 39:85–105. [aMO]
Manktelow, K. I., Sutherland, E. J. & Over, D. E. (1995) Probabilistic factors in deontic reasoning. *Thinking and Reasoning* 1:201–20. [aMO]
Manning, C. & Schütze, H. (1999) *Foundations of statistical natural language processing*. MIT Press. [rMO]
Markovits, H. & Barrouillet, P. (2002) The development of conditional reasoning: A mental model account. *Developmental Review* 22:5–36. [KO, NS]
Marr, D. (1982) *Vision: A computational investigation into the human representation and processing of visual information*. Freeman. [ELK, CRMM, JDN, rMO]
Martignon, L. & Blackmond-Laskey, K. (1999) Bayesian benchmarks for fast and frugal heuristics. In: *Simple heuristics that make us smart*, ed. G. Gigerenzer, P. M. Todd & the ABC Research Group, pp. 169–88. Oxford University Press. [rMO]
McCarthy, J. & Hayes, P. J. (1969) Some philosophical problems from the standpoint of artificial intelligence. In: *Machine intelligence, vol. 4*, ed. B. Meltzer & D. Michie. Edinburgh University Press. [aMO]
McClelland, J. L. (1998) Connectionist models and Bayesian inference. In: *Rational models of cognition*, ed. M. Oaksford & N. Chater, pp. 21–53. Oxford University Press. [aMO]
McKenzie, C. R. M. (2003) Rational models as theories – not standards – of behavior. *Trends in Cognitive Sciences* 7(9):403–406. [JDN]
(2004) Framing effects in inference tasks – and why they are normatively defensible. *Memory and Cognition* 32:874–85. [CRMM]
McKenzie, C. R. M. & Amin, M. B. (2002) When wrong predictions provide more support than right ones. *Psychonomic Bulletin and Review* 9:821–28. [CRMM]
McKenzie, C. R. M., Ferreira, V. S., Mikkelsen, L. A., McDermott, K. J. & Skrable, R. P. (2001) Do conditional statements target rare events? *Organizational Behavior and Human Decision Processes* 85:291–309. [CRMM, arMO]
McKenzie, C. R. M. & Mikkelsen, L. A. (2000) The psychological side of Hempel's paradox of confirmation. *Psychonomic Bulletin and Review* 7:360–66. [CRMM, arMO]
(2007) A Bayesian view of covariation assessment. *Cognitive Psychology* 54:33–61. [CRMM, arMO]
McKenzie, C. R. M. & Nelson, J. D. (2003) What a speaker's choice of frame reveals: Reference points, frame selection, and framing effects. *Psychonomic Bulletin and Review* 10:596–602. [CRMM]
Milch, B., Marthi, B. & Russell, S. (2004) BLOG: Relational modeling with unknown objects. In: *ICML 2004 workshop on statistical relational learning and its connections to other fields, Banff, Alberta, Canada*, ed. T. Dietterich, L. Getoor & K. Murphy, pp. 67–73. Available at: www.cs.umd.edu/projects/srl2004/srl2004_complete.pdf [TLG]
Milne, P. (1995) A Bayesian defence of Popperian science? *Analysis* 55:213–15. [rMO]
(1996) log[P(h|eb)/P(h|b)] is the one true measure of confirmation. *Philosophy of Science* 63:21–26. [rMO]
Moore, G. E. (1903) *Principia ethica*. Cambridge University Press. [rMO]
Murphree, W. A. (1991) *Numerically exceptive logic: A reduction of the classical syllogism.* Peter Lang. [NP]
Navarro, D. J., Griffiths, T. L., Steyvers, M. & Lee, M. D. (2006) Modeling individual differences using Dirichlet processes. *Journal of Mathematical Psychology* 50:101–22. [rMO, NS]
Nelson, J. D. (2005) Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review* 112(4):979–99. [JDN, arMO]
(2008) Towards a rational theory of human information acquisition. In: *The probabilistic mind: Prospects for rational models of cognition*, ed. M. Oaksford & N. Chater, pp. 143–63. Oxford University Press. [JDN]
Nelson, J. D., McKenzie, C. R. M., Cottrell, G. W. & Sejnowski, T. J. (submitted) Experience matters: Information acquisition optimizes probability gain. [JDN]
Newell, A., Shaw, J. C. & Simon, H. A. (1958) Chess-playing programs and the problem of complexity. *IBM Journal of Research and Development* 2:320–25. [aMO]
Newell, A. & Simon, H. A. (1972) *Human problem solving*. Prentice-Hall. [aMO]
Newstead, S. E., Handley, S. J. & Buck, E. (1999) Falsifying mental models: Testing the predictions of theories of syllogistic reasoning. *Memory and Cognition* 27:344–54. [aMO]
Nickerson, R. S. (1996) Hempel's paradox and Wason's selection task: Logical and psychological puzzles of confirmation. *Thinking and Reasoning* 2:1–32. [JDN, aMO]
Novick, L. R. & Cheng, P. W. (2004) Assessing interactive causal influence. *Psychological Review* 111:455–85. [aMO]
Oaksford, M. (2004a) Conditional inference and constraint satisfaction: Reconciling probabilistic and mental models approaches? Paper presented at the 5th International Conference on Thinking, July 22–24, 2004. University of Leuven, Leuven, Belgium. [aMO]
(2004b) Reasoning. In: *Cognitive psychology*, ed. N. Braisby & A. Gellatly, pp. 418–55. Oxford University Press. [aMO]
Oaksford, M. & Brown, G. D. A., eds. (1994) *Neurodynamics and psychology.* Academic Press. [UH]
Oaksford, M. & Chater, N. (1991) Against logicist cognitive science. *Mind and Language* 6:1–38. [arMO]

(1994) A rational analysis of the selection task as optimal data selection. *Psychological Review* 101:608–31. [JDN, arMO, UH]
(1996) Rational explanation of the selection task. *Psychological Review* 103:381–91. [aMO]
(1998a) *Rationality in an uncertain world*. Psychology Press. [aMO, WS]
eds. (1998b) *Rational models of cognition.* Oxford University Press. [UH, arMO]
(2001) The probabilistic approach to human reasoning. *Trends in Cognitive Sciences* 5:349–57. [E-JW]
(2002) Common sense reasoning, logic and human rationality. In: *Common sense, reasoning and rationality*, ed. R. Elio, pp. 174–214. Oxford University Press. [rMO]
(2003a) Conditional probability and the cognitive science of conditional reasoning. *Mind and Language* 18:359–79. [aMO, DPO]
(2003b) Optimal data selection: Revision, review, and reevaluation. *Psychonomic Bulletin and Review* 10:289–318. [aMO, E-JW]
(2007) *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press. [NA, HB, DD, WDN, JStBTE, TLG, UH, GSH, ELK, I-mL, CRMM, JDN, KO, arMO, DPO, DEO, NP, FHP, GP. KS, NS, WS, E-JW]
(2008) Probability logic and the *Modus Ponens–Modus Tollens* asymmetry in conditional inference. In: *The probabilistic mind*: *Prospects for Bayesian cognitive science*, ed. N. Chater & M. Oaksford, pp. 97–120. Oxford University Press. [arMO]
(in press) Conditionals and constraint satisfaction: Reconciling mental models and probabilistic approaches? In: *Cognition and conditionals*: *Probability and logic in human thought*, ed. M. Oaksford & N. Chater. Oxford University Press. [rMO]
Oaksford, M., Chater, N. & Grainger, B. (1999) Probabilistic effects in data selection. *Thinking and Reasoning* 5:193–244. [arMO]
Oaksford, M., Chater, N., Grainger, B. & Larkin, J. (1997) Optimal data selection in the reduced array selection task (RAST). *Journal of Experimental Psychology: Learning, Memory and Cognition*, 23:441–58. [aMO]
Oaksford, M., Chater, N. & Larkin, J. (2000) Probabilities and polarity biases in conditional inference. *Journal of Experimental Psychology: Learning, Memory and Cognition* 26:883–89. [KO, arMO]
Oaksford, M. & Hahn, U. (2007) Induction, deduction and argument strength in human reasoning and argumentation. In: *Inductive reasoning*, ed. A. Feeney & E. Heit. pp. 269–301. Cambridge University Press. [rMO]
Oaksford, M. & Moussakowski, M. (2004) Negations and natural sampling in data selection: Ecological vs. heuristic explanations of matching bias. *Memory and Cognition* 32:570–81. [arMO]
Oaksford, M., Roberts, L. & Chater, N. (2002) Relative informativeness of quantifiers used in syllogistic reasoning. *Memory and Cognition* 30:138–49. [aMO]
Oaksford, M. & Stenning, K. (1992) Reasoning with conditionals containing negated constituents. *Journal of Experimental Psychology: Learning, Memory and Cognition* 18:835–54. [aMO]
Oaksford, M. & Wakefield, M. (2003) Data selection and natural sampling: Probabilities do matter. *Memory and Cognition* 31:143–54. [arMO]
Oberauer, K. (2006) Reasoning with conditionals: A test of formal models of four theories. *Cognitive Psychology* 53:238–83. [UH, KO, rMO]
Oberauer, K., Weidenfeld, A. & Fischer, K. (2007) What makes us believe a conditional? The roles of covariation and causality. *Thinking and Reasoning* 13:340–69. [KO]
Oberauer, K., Weidenfeld, A. & Hörnig, R. (2004) Logical reasoning and probabilities: A comprehensive test of Oaksford and Chater (2001) *Psychonomic Bulletin and Review* 11:521–27. [KO, arMO]
Oberauer, K. & Wilhelm, O. (2003) The meaning(s) of conditionals: Conditional probabilities, mental models and personal utilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29:680–93. [I-mL, aMO]
Oberauer, K., Wilhelm, O. & Diaz, R. R. (1999) Bayesian rationality for the Wason selection task? A test of optimal data selection theory. *Thinking and Reasoning* 5:115–44. [aMO, E-JW]
O'Brien, D. P. (1993) Mental logic and irrationality: We can put a man on the moon, so why can't we solve those logical reasoning problems? In: *Rationality: Psychological and philosophical perspectives*, ed. K. I. Manktelow & D. E. Over. Routledge. [DPO]
(2004) Mental-logic theory: What it proposes, and reasons to take this proposal seriously. In: *The nature of reasoning*, ed. J. Leighton & R. J. Sternberg, pp. 205–33. Cambridge University Press. [DPO]
O'Brien, D. P. & Manfrinati, A. (in press) The mental logic theory of conditional propositions. In: *The psychology of conditionals*, ed. M. Oaksford & N. Chater. Oxford University Press. [DPO]
Osherson, D. N., Stob, M. & Weinstein, S. (1988) Mechanical learners pay a price for Bayesianism. *Journal of Symbolic Logic* 53:1245–51. [DD]
Over, D. E. & Evans, J. St. B. T (1994) Hits and misses: Kirby on the selection task. *Cognition* 52:235–43. [aMO]
Over, D. E., Hadjichristidis, C., Evans, J. St. B. T., Handley, S. J. & Sloman, S. A. (2007) The probability of causal conditionals. *Cognitive Psychology* 54:62–97. [I-mL, DEO, rMO]
Over, D. E. & Jessop, A. (1998) Rational analysis of causal conditionals and the selection task. In: *Rational models of cognition*, ed. M. Oaksford & N. Chater, pp. 399–414. Oxford University Press. [aMO]
Pearl, J. (1988) *Probabilistic reasoning in intelligent systems*. Morgan Kaufmann. [arMO]
(2000) *Causality: Models, reasoning and inference*. Cambridge University Press. [arMO]
Perham, N. & Oaksford, M. (2005) Deontic reasoning with emotional content: Evolutionary psychology or decision theory? *Cognitive Science* 29:681–718. [aMO]
Peters, S. & Westerståhl, D. (2006) *Quantifiers in language and logic*. Oxford University Press. [NP]
Peterson, P. L. (2000) *Intermediate quantifiers. Logic, linguistics, and Aristotelian semantics*. Ashgate. [NP]
Pfeifer, N. & Kleiter, G. D. (2005a) Coherence and nonmonotonicity in human reasoning. *Synthese* 146(1–2):93–109. [GP, NP]
(2005b) Towards a mental probability logic. *Psychologica Belgica* 45(1):71–99. Updated version available at: http://www.users.sbg.ac.at/~pfeifern/problog.htm [NP]
(2006) Inference in conditional probability logic. *Kybernetika* 42:391–404. [NP]
(in press a) Framing human inference by coherence based probability logic. *Journal of Applied Logics*. [NP]
(in press b) The conditional in mental probability logic. In: *Cognition and conditionals: Probability and logic in human thoughts*, ed. M. Oaksford & N. Chater. Oxford University Press. [NP]
Piaget, J. (1957) *Logic and psychology*. Basic Books. [GSH]
Piaget, J. & Inhelder, B. (1975) *The origin of the idea of chance in children*. W. W. Norton. [NS]
Pitt, M. A., Myung, I. J. & Zhang, S. (2002) Toward a method of selecting among computational models of cognition. *Psychological Review* 109:472–91. [HB]
Poletiek, F. H. (1996) Paradoxes of falsification. *Quarterly Journal of Experimental Psychology* 49A:447–62. [FHP]
(2001) *Hypothesis-testing behavior*. Psychology Press. [FHP]
Poletiek, F. H. & Berndsen, M. (2000) Hypothesis testing as risk behavior with regard to beliefs. *Journal of Behavioral Decision Making* 13:107–23. [FHP]
Politzer, G. & Braine, M. D. (1991) Responses to inconsistent premises cannot count as suppression of valid inferences. *Cognition* 38:103–08. [aMO]
Popper, K. R. (1935/1959) *The logic of scientific discovery*. Basic Books. [aMO]
(1963/1978) *Conjectures and refutations*, 4th edition. Routledge and Kegan Paul. [FHP]
Putnam, H. (1974) The "corroboration" of theories. In: *The philosophy of Karl Popper, vol. 2*, ed. A. Schilpp. Open Court. [aMO]
Pylyshyn, Z., ed. (1987) *The robot's dilemma: The frame problem in artificial intelligence.* Ablex. [aMO]
Quine, W. V. O. (1953) *From a logical point of view.* Harvard University Press. [aMO]
Raby, C. R., Alexis, D. M., Dickinson, A. & Clayton, N. S. (2007) Planning for the future by Western scrub-jays. *Nature* 445:919–21. [ELK]
Ramsey, F. P. (1931/1990a) General propositions and causality. In: *Philosophical papers*, ed. D. H. Mellor, pp. 145–63. Cambridge University Press. [I-mL]
(1931/1990b) *The foundations of mathematics and other logical essays.* Routledge and Kegan Paul. [aMO]
Rao, R. P. N., Olshausen, B. A. & Lewicki, M. S., eds. (2002) *Probabilistic models of the brain: Perception and neural function*. MIT Press. [rMO]
Reichenbach, H. (1938) *Experience and prediction.* University of Chicago Press. [TLG]
Reiter, R. (1980) A logic for default reasoning, *Artificial Intelligence* 13:81–132. [aMO]
Restall, G. (1996) Information flow and relevant logics. In: *Logic, language and computation*, ed. J. Seligman, & D. Westerståhl, pp. 463–78 CSLI Publications. [rMO]
Rips, L. J. (1983) Cognitive processes in propositional reasoning. *Psychological Review* 90:38–71. [aMO]
(1994) *The psychology of proof.* MIT Press. [arMO]
(2002) Reasoning imperialism. In: *Common sense, reasoning, and rationality*, ed. R. Elio, pp. 215–35. Oxford University Press. [WS]
Rissanen, J. J. (1989) *Stochastic complexity and statistical inquiry*. World Scientific. [rMO]
Rosch, E. (1975) Cognitive representation of semantic categories. *Journal of experimental psychology: General* 104:192–233. [aMO]
Rosser, J. B. (1939) An informal exposition of proofs of Gödel's theorem and Church's theorem. *Journal of Symbolic Logic* 4:53–60. [DPO]

Sanborn, A. N., Griffiths, T. L. & Navarro, D. J. (2006) A more rational model of categorization. In: *Proceedings of the 28th Annual Conference of the Cognitive Science Society.* Erlbaum. [TLG]

Savage, L. J. (1972) *Foundations of statistics*. Dover. [DD]

Schroyens, W. (in press) Logic and/in psychology: The paradoxes of material implication and psychologism in the cognitive science of human reasoning. In: *Cognition and conditionals: Probability and logic in human thoughts*, eds. M. Oaksford & N. Chater. Oxford University Press. [WS]

Schroyens, W. & Schaeken, W. (2003) A critique of Oaksford, Chater and Larkin's (2000) conditional probability model of conditional reasoning. *Journal of Experimental Psychology: Learning, Memory and Cognition* 29:140–49. [UH, aMO]

Schulte, O. (1999) The logic of reliable and efficient inquiry. *Journal of Philosophical Logic* 28:399–438. [DD]

Shafer, G. (1976) *A mathematical theory of evidence*. Princeton University Press. [rMO, GP]

Shannon, C. E. & Weaver, W. (1949) *The mathematical theory of communication*. University of Illinois Press. [rMO]

Shepard, R. N. (1987) Towards a universal law of generalization for psychological science. *Science* 237:1317–23. [TLG]

(1995) Mental universals: Toward a twenty-first century science of mind. In: *The science of the mind: 2001 and beyond*, ed. R. L. Solso & D. W. Massaro, pp. 50–62. Oxford University Press. [TLG]

Sher, S. & McKenzie, C. R. M. (2006) Information leakage from logically equivalent frames. *Cognition* 101:467–94. [CRMM, rMO]

(2008) Framing effects and rationality. In: *The probabilistic mind: Prospects for Bayesian cognitive science*, ed. N. Chater & M. Oaksford, pp. 79–96. Oxford University Press. [CRMM]

Shi, L., Feldman, N. & Griffiths, T. L. (2008) Performing Bayesian inference with exemplar models. In: *Proceedings of the Thirtieth Annual Conference of the Cognitive Science Society*, ed. B. Love, K. McRae & V. Sloutsky, pp. 745–50. Cognitive Science Society. [TLG]

Shultz, T. R. (2007) The Bayesian revolution approaches psychological development. *Developmental Science* 10:357–64. [NS]

Skov, R. B. & Sherman, S. J. (1986) Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology* 22:93–121. [JDN]

Sloman, S. A. (1996) The empirical case for two systems of reasoning. *Psychological Bulletin* 119:3–22. [rMO]

(2005) *Causal models*. Oxford University Press. [rMO]

Slowiaczek, L. M., Klayman, J., Sherman, S. J. & Skov, R. B. (1992) Information selection and use in hypothesis testing: What is a good question, and what is a good answer? *Memory and Cognition* 20:392–405. [JDN]

Smolensky, P. (1990) Tensor product variable binding and the representation of symbolic structures in connectionist networks. *Artificial Intelligence* 46:159–216. [UH]

Sobel, D. M., Tenenbaum, J. B. & Gopnik, A. (2004) Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science* 28:303–33. [rMO, NS]

Sobel, J. H. (2004) Probable modus ponens and modus tollens and updating on uncertain evidence. Unpublished manuscript, Department of Philosophy, University of Toronto, Scarborough. Available at: www.scar.toronto.ca/~sobel/Conf/Disconf.pdf. [rMO]

Sober, E. (2002) Intelligent design and probability reasoning. *International Journal for Philosophy of Religion* 52:65–80. [arMO]

Stanovich, K. E. (1999) *Who is rational? Studies of individual differences in reasoning*. Elrbaum. [JStBTE]

(2008) Individual differences in reasoning and the algorithmic/intentional level distinction in cognitive science. In: *Reasoning: Studies of human inference and its foundations*, ed. L. Rips & J. Adler, pp. 414–36. Cambridge University Press. [rMO]

Stanovich, K. E. & West, R. F. (2000) Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences* 23:645–65. [rMO]

Stenning, K. & van Lambalgen, M. (2005) Semantic interpretation as reasoning in non-monotonic logic: The real meaning of the suppression task. *Cognitive Science* 29(6):919–60. [KS]

(2008a) *Human reasoning and cognitive science*. MIT Press. [KS]

(2008b) Logic in a noisy world. In: *Cognition and conditionals: Probability and logic in human thought*, eds. M. Oaksford & N. Chater. Oxford University Press. [KS]

Sternberg, R. J. (1999) Intelligence as developing expertise. *Contemporary Educational Psychology* 24:359–75. [NS]

Stevenson, R. J. & Over, D. E. (1995) Deduction from uncertain premises. *The Quarterly Journal of Experimental Psychology* 48A:613–43. [DEO, rMO]

(2001) Reasoning from uncertain premises. Effects of expertise and conversational context. *Thinking and Reasoning* 7:367–90. [DEO]

Stewart, N., Chater, N. & Brown, G. D. A. (2006) Decision by sampling. *Cognitive Psychology* 53:1–26. [arMO]

Taplin, J. E. (1971) Reasoning with conditional sentences. *Journal of Verbal Learning and Verbal Behavior* 10:219–25. [I-mL]

Taplin, J. E. & Staudenmayer, H. (1973) Interpretation of abstract conditional sentences in deductive reasoning. *Journal of Verbal Learning and Verbal Behavior* 12:530–42. [I-mL]

Teller, P. (1973) Conditionalization and observation. *Synthese* 26:218–38. [DD]

Tenenbaum, J. B. (1999) A Bayesian framework for concept learning. Doctoral dissertation, Brain and Cognitive Sciences Department, MIT. [rMO]

Tenenbaum, J. B. & Griffths, T. L. (2001) Structure learning in human causal induction. In: *Advances in neural information processing systems, vol. 13*, ed. T. K. Keen, T. G. Dietterich & V. Tresp, pp. 59–65. MIT Press. [rMO]

Tenenbaum, J. B., Griffiths, T. L. & Kemp, C. (2006) Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Science* 10:309–18. [TLG]

Tenenbaum, J. B., Kemp, C. & Shafto, P. (2007) Theory based Bayesian models of inductive reasoning. In: *Inducive reasoning*, ed. A. Feeney & E. Heit, pp. 167–204. Oxford University Press. [UH, rMO]

Thaler, R. H. (2005) *Advances in behavioral finance, Vol. II.* Princeton University Press. [rMO]

Todd, P. M. & Gigerenzer, G. (2003) Bounding rationality to the world. *Journal of Economic Psychology* 24:143–65. [HB]

Toussaint, M., Harmeling, S. & Storkey, A. (2006) Probabilistic inference for solving (PO)MDPs. Technical Report EDI-INF-RR-0934, University of Edinburgh. [rMO]

van Fraassen, B. C. (1980) *The scientific image*. Oxford University Press. [ELK]

Verhaeghen, P. & Salthouse, T. A. (1997) Meta-analyses of age-cognition relations in adulthood: Estimates of linear and nonlinear age effects and structural models. *Psychological Bulletin* 122:231–49. [NS]

Vermeij, G. J. (2004) *Nature: An economic history.* Princeton University Press. [ELK]

Viswanathan, G. M., Buldyrev, S. V., Havlin, S., da Luz, M. G., Raposo, E. P. & Stanley, H. E. (1999) Optimizing the success of random searches. *Nature* 401:911–14. [JDN]

Vitányi, P. M. B. & Li, M. (2000) Minimum description length induction, Bayesianism, and Kolmogorov complexity. *IEEE Transactions on Information Theory* IT-46446–64. [rMO]

von Helmholtz, H.(1910/1925) *Physiological optics. Volume III. The theory of the perception of vision.* Dover. (Translated from 3rd German edition, 1910). [rMO]

Wagner, C. G. (2004) Modus tollens probabilized. *British Journal for Philosophy of Science* 55:747–53. [aMO]

Walton, D. N. (1989) *Informal logic*. Cambridge University Press. [rMO]

Wang, Z.-J. (1999) A two-component model of conditional reasoning. Unpublished Master's Thesis, National Chung-Cheng University, Chia-Yi, Taiwan. [I-mL]

Wells, G. L. & Lindsay, R. C. L. (1980) On estimating the diagnosticity of eyewitness nonidentifications. *Psychological Bulletin* 88:776–84. [JDN]

Williamson, J. & Gabbay, D., eds. (2003) Special issue on combining logic and probability. *Journal of Applied Logic* 1:135–308. [rMO]

Yama, H. (2001) Matching versus optimal data selection in the Wason selection task. *Thinking and Reasoning* 7:295–311. [aMO]

Yuille, A. L. & Kersten, D. (2006) Vision as Bayesian Inference: Analysis by Synthesis? *Trends in Cognitive Sciences* 10:301–08. [aMO]

Zadeh, L. A. (1975) Fuzzy logic and approximate reasoning. *Synthese* 30:407–28. [rMO]

Zeelenberg, M., Van Dijk, W. W., Manstead, A. S. R. & Van der Pligt, J. (2000) On bad decisions and disconfirmed expectancies: The psychology of regret and disappointment. *Cognition and Emotion* 14:521–41. [aMO]