ANNUAL REVIEWS

Annual Review of Psychology

New Paradigms in the Psychology of Reasoning

Mike Oaksford[1] and Nick Chater[2]

[1]Department of Psychological Sciences, Birkbeck, University of London, London WC1E 7HX, United Kingdom; email: mike.oaksford@bbk.ac.uk

[2]Nick Chater, Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom; email: Nick.Chater@wbs.ac.uk

Annu. Rev. Psychol. 2020. 71:305–30

First published as a Review in Advance on September 12, 2019

The *Annual Review of Psychology* is online at psych.annualreviews.org

https://doi.org/10.1146/annurev-psych-010419-051132

Keywords

reasoning, inference, persuasion, argument, probability, Bayes' theorem, logic, world knowledge, nonmonotonicity, fallacies

Abstract

The psychology of verbal reasoning initially compared performance with classical logic. In the last 25 years, a new paradigm has arisen, which focuses on knowledge-rich reasoning for communication and persuasion and is typically modeled using Bayesian probability theory rather than logic. This paradigm provides a new perspective on argumentation, explaining the rational persuasiveness of arguments that are logical fallacies. It also helps explain how and why people stray from logic when given deductive reasoning tasks. What appear to be erroneous responses, when compared against logic, often turn out to be rationally justified when seen in the richer rational framework of the new paradigm. Moreover, the same approach extends naturally to inductive reasoning tasks, in which people extrapolate beyond the data they are given and logic does not readily apply. We outline links between social and individual reasoning and set recent developments in the psychology of reasoning in the wider context of Bayesian cognitive science.

Contents

INTRODUCTION

Psychologists have studied adult verbal reasoning for more than half a century, initially viewing binary logic, in which propositions can only be true or false, as the appropriate standard of correct inference. Typical experiments focus on a small range of verbal inferences using the logical terms of natural language such as *and*, *or*, *not*, *if…then*, *all*, *some*. Participants might be given some premises—for example, *if the key is turned*, *the car starts*, and *the key is turned*—and then asked whether they endorse the conclusion that *the car starts*. Here, if they are reasoning using classical logic, they should endorse the conclusion. Such narrow logical reasoning tasks gain wider significance if thought, in general, is viewed as logical inference over beliefs encoded in a language of thought [as in symbolic artificial intelligence (AI) and classical cognitive science; see Fodor 1975].

Over the last 25 years, what Over (2009) has called the new paradigm in reasoning has emerged, a shift especially associated with switching from binary logic to Bayesian probability theory as the standard of good inference (Oaksford & Chater 1994). Some have interpreted this shift as merely swapping one normative, equally logical, theory of reasoning for another (Elqayam & Evans 2011). Yet the new paradigm embodies two other fundamental changes: (*a*) It focuses on knowledge-rich inference based on content and background knowledge, not just structural reasoning based on logical form, and (*b*) it prioritizes the social role of reasoning in communication, argument, and persuasion (Hahn & Oaksford 2007; Mercier & Sperber 2011, 2017; Oaksford & Hall 2016). In this article, we review argumentation before individual reasoning, despite its later emergence in the new paradigm.

We begin with an overview of the transition to the new paradigm in the psychology of reasoning in the 1990s, introducing the main themes in its later development. The next sections focus in turn on social and individual reasoning and consider how these interact. We finish by setting these developments in the context of broader transformations in the cognitive and brain sciences, before describing challenges and future directions. This review is by necessity highly selective, excluding many important topics (e.g., dual processes, meta-reasoning, and causal and counterfactual reasoning) and focusing on what we see as the central themes of the new paradigm and its place in the cognitive sciences.

NEW PARADIGMS FOR REASONING

New paradigms emerge in reaction to previous paradigms. Hence, we begin by introducing the logical paradigm implicit in the early psychology of reasoning and still active today.

The Logical Paradigm

The logical paradigm uses classical binary logic as its normative theory. For example, consider the conditional (if p then q) inference relating antecedent, p, to consequent, q.

Proposition 1.

	If I turn the key, the car starts.	If p then q	(A)
	The key is turned.	p	(B)
Therefore	The car starts.	q	(C) A & B, by MP

In binary logic, the conclusion, *the car starts*, follows by the logical law of modus ponens (MP). This inference is logically valid: That is, if the premises A and B are true, so must be the conclusion (C). Such inferences are converted into experiments using what Evans (2017; see also George 1995) calls the assumption-based paradigm. Participants are instructed to assume the premises are true, even if they do not believe them, and to draw the inference with those, and only those, assumptions (we shall see later that these instructions may be highly unnatural for participants).

The two most prominent old-paradigm theories make different assumptions about mental representations and processes. Mental logic (e.g., Rips 1994) proposes that the mind mirrors syntactic representations and the application of logical rules, as above. By contrast, mental models theory (MMT) (Johnson-Laird 1983, Johnson-Laird & Byrne 2002) assumes that people reason over representations that attempt to capture what the premises mean. For propositional reasoning, a mental model is equivalent to what logicians call the disjunctive normal form (DNF) of a complex proposition (one that relates more than one variable, e.g., p and q). For example, the DNF of a conditional *if p then q* is (p and q) or ($\neg p$ and q) or ($\neg p$ and $\neg q$), where $\neg$ = not. In MMT, the inference in Proposition 1 follows because, in the only disjunct in which p is true, q is also true.

Problems for the old paradigm. Problems for old-paradigm approaches emerge if we consider our conditional (A) more carefully. If we are told that *the car did not start*, another logical law, modus tollens (MT), invites us to infer that *I did not turn the key*. However, if the car not starting were mentioned in conversation, the opposite conclusion would be drawn: that I did turn the key, but the car failed to start (Oaksford & Chater 2007, 2013). This implication is based both on our understanding of relevant aspects of the world (cars do not start spontaneously, they are mostly immobile, etc.) and on the norms of conversation (Sperber & Wilson 1986). We infer the existence of a failed key-turning attempt; otherwise, the car not starting would not be worth mentioning.

Note, though, that listeners thus spontaneously infer a counterexample to the very rule that is being implicitly applied. From a conventional logical standpoint, this would involve rejecting the conditional as false and no longer using it to guide our future actions. However, the very statement that *the car did not start this morning* only gains significance on the assumption that the rule is generally reliable: Cars normally start when their keys are turned, but this one did not. So, even in the simplest conditional inference, our reasoning depends not on pure principles of logic, but on a combination of background knowledge and conversational principles.

This example also highlights a violation of a fundamental logical law presumed to govern conditionals: strengthening of the antecedent, which proposes that if *if p then q* is true, then *if p and* *r then q* must also be true. However, *if I turn the key, the car starts* certainly does not imply that *if I turn the key and the gas tank is empty, the car starts.* This ubiquitous pattern in everyday inference is known as nonmonotonicity, which raises the specter of computational intractability: nonmonotonic reasoning in AI often requires explosive amounts of computation (e.g., Oaksford & Chater 1991, 2007).

Five fundamental problems. There are, we suggest, five fundamental problems for the old paradigm:

1. Logical reasoning is knowledge poor, depending only on the logical forms of the premises. However, this leads to endorsing bizarre inferences, such as from *the moon is not made of cheese* to *if the moon is made of cheese, then life exists on other planets*.
2. Any background knowledge (added as additional premises) must be consistent, on pain of inferential explosion: In classical logic, and many other logics, any proposition follows from a contradiction by the logical law of *ex falso quodlibet*. However, consistency is very difficult to achieve, let alone verify.
3. There can be no learning. If inductive learning were based on logical entailment (Carnap 1945), "it would in fact be impossible to learn from experience" (Salmon 1967, pp. 731–32).
4. Logical reasoning is monotonic, so how is people's normally nonmonotonic inference to be captured rationally?
5. The old paradigm concentrates on individual reasoning. However, as in our car-starting example, inference often makes sense only in a conversational context.

Each of these problems is cast in a new light if we see reasoning as primarily social: as a process of argumentation and persuasion, which is prior to understanding individual reasoning (Hahn & Oaksford 2007, Mercier & Sperber 2011). Moreover, as we shall see, the new paradigm may also help resolve the problems we have just raised.

The Probabilistic Paradigm

A probabilistic Bayesian viewpoint sees belief as inherently uncertain and reasoning as concerned with updating our uncertain beliefs. Thus, we can rarely be certain of the truth of the premises over which we reason. Moreover, the function of evidence and argument is to change our belief in a conclusion, and such conclusions might form the basis for action (Harman 1986). Degrees of belief can be captured by probabilities (this is the subjective, or Bayesian, interpretation of probability); probability theory is the mathematical theory of how to reason with degrees of belief. The general form for the probability of a proposition, p, is $\Pr(p|B)$, that is, any probability assignment is conditional on background knowledge B (though we omit the $|B$ below). Consequently, the new paradigm is knowledge rich from the outset (problem 1 above).

The new paradigm developed from two research strands in the mid-1990s. One strand originated in Anderson's (1991) rational analysis program, which attempts to describe human behavior as optimally adapted to the environment. The environment is captured by prior probabilities; new information updates these priors using Bayes' theorem, an approach closely aligned with Bayesian philosophy of science (e.g., Howson & Urbach 1989). Bayes' theorem provides a normative account of learning (problem 3). In the psychology of reasoning, this approach was first applied to confirmation bias in Wason's selection task (Oaksford & Chater 1994), an experimental analog of testing scientific hypotheses that we will discuss later. The second strand originated from philosophical concerns about the adequacy of classical binary logic to characterize human reasoning, given the uncertainty of the real world (Edgington 1995). Related concerns arose because classic logic is monotonic, whereas everyday reasoning is nonmonotonic (Oaksford & Chater 1991). Empirical studies confirmed that premise uncertainty (e.g., Do cars invariably start when one turns the key? Was the key really turned?) had predictable effects on conclusion evaluation (e.g., George 1995, Stevenson & Over 1995). Moreover, work on causal conditional reasoning (Cummins et al. 1991) had already demonstrated examples of nonmonotonicity called suppression effects (Byrne 1989).

These two strands converged on the interpretation of the most important logical term in natural language, *if…then*: Both adopted what Edgington (1995) christened the Equation: that the probability of the conditional Pr(*if p then q*) is the conditional probability $\Pr(q|p)$. Logicians call such a conditional a probability conditional. Strengthening the antecedent is not valid for the probability conditional [i.e., $\Pr(q|p\&r)$ can be less than $\Pr(q|p)$], leading to nonmonotonicity not allowed in classical logic (problem 4). The Equation also permitted the generalization of the new paradigm to quantified syllogistic reasoning, involving the logical terms *all*, *some*, *none*, and *some…not* (Chater & Oaksford 1999), and to so-called polarity biases created by negations in conditional sentences (Oaksford et al. 2000). Moreover, several studies have confirmed the effects of premise uncertainty on conditional reasoning (Politzer & Carles 2001, Stevenson & Over 2001) and provided direct empirical evidence for the Equation (e.g., Evans et al. 2003, Oberauer & Wilhelm 2003).

The generalization to quantified syllogistic reasoning (Chater & Oaksford 1999) embodied many recent developments, defining a notion of (*a*) probabilistic or *p*-validity (Adams 1998) that relied on (*b*) deducing a coherence interval for the conclusion (Coletti & Scozzafava 2002), assuming that the premises defined (*c*) a dependency graph (Pearl 1988). All of these elements feature in the later development of the new paradigm.

In the logical paradigm, the natural extension is to further logical terms, retaining the focus on individual and intrapersonal reasoning (for example, Hinterecker et al. 2016). By contrast, Hahn and Oaksford (Hahn & Oaksford 2007, Oaksford & Hahn 2004) generalized the new paradigm to argumentation, setting reasoning in its broader social context (problem 5; see Mercier & Sperber 2011, 2017).

The Social Character of Reasoning

The brain has only shallow and fragmentary knowledge to deal with a continual stream of new situations and urgent decisions (Oaksford & Chater 2007, Rozenblit & Keil 2002). Reasoning begins from fragmented knowledge: that cars normally start when we turn the ignition key, that they typically need gas, that our car did not start as expected this morning, and so on. Normative theories of reasoning, such as logic and probability theory, can spell out under what conditions there are contradictions between these fragments—for example, some fragments imply that when we turn the key, the car starts; others tell us that it does not. Normative theories tell us something about what counts as resolving a contradiction—for example, realizing the car had no gas reconciles the general rule and the specific counterexample. However, these theories cannot tell us which of the many possible ways of reconciling contradictions should be preferred (e.g., we could conclude that the key was not turned or that the general rule is false).

The gaps and inconsistencies in our individual knowledge are partially ameliorated by the social distribution of knowledge (Sloman & Fernbach 2017) and the exchange of information in communication (Oaksford & Hall 2016). For example, if the car does not start, we will often not know why, or how to fix it; we will call a local mechanic. More generally, we do not individually need to repair all the gaps and contradictions in our fragmentary knowledge; we understand and cope with the world collectively (addressing problem 2).

The communicative aspect of reasoning is also embodied in the very interpretation of the statements over which we reason (Hilton 1995). Suppose you say *I turned the key, but the car did not start*. There is no contradiction with background knowledge unless the key is the ignition key of your car, was turned when in the ignition, and in the right direction. Both speaker and listener must do substantial inferential work to agree on interpretations of the relevant linguistic terms. Sometimes the result of an argument is to uncover misaligned interpretation (e.g., "Oh I see, that was the key to the other car!"). In any case, verbal reasoning only begins when we have settled issues of interpretation (Stenning & Cox 2006).

Language facilitates argumentation through which we can resolve disagreements by filling in our fragmentary and shallow knowledge. Argumentation involves persuading an audience of a disputed claim by offering a series of rational moves that support or refute it (Van Eemeren & Grootendorst 2004). Argumentation is knowledge rich: Whether new evidence or reasoning is convincing depends on what we already know. Moreover, argumentation is local and dynamic: Interlocutors are attempting to win a local game by repelling specific counterarguments. So the focus is on the local disagreement, in the light of global background knowledge. Language and argumentation allow isolated individuals to reduce their ignorance and resolve disagreements.

Language and argumentation may have culturally coevolved in a way that can hint at the origins of our norms of reasoning. As hunter-gatherer groups grew larger and more complex, they needed to forage more widely and to bring back information about the location of food sources (Tomasello 2014). Of course, the waggle dance of honeybees communicates similar information about the location of food to the group (hive), which responds reflexively by sending more bees to the location. However, in a hominid group, other members may query the information conveyed by the returning foragers [Bennett 1989 (1964)] or provide counterarguments (e.g., "We have been there and found nothing," or "There was danger"). This response is what Bennett [1989 (1964)] referred to as reasoned denial. Resolving what to do provides the key ingredients for rational debate: understanding negation ("No, we should not go there"), providing reasons ("There is danger"), and, in particular, acknowledging the principle of noncontradiction [one cannot believe both that we should and should not go to this location, i.e., $\neg(p \text{ and } \neg p)$]. Adherence to this principle is implicit in initiating the disagreement that subsequent discussion must resolve; and, as Millikan (2006, p. 125) argued, this "is a form of rationality that it seems less likely that nonhuman animals can achieve." For Aristotle, the principle of noncontradiction was a self-evident axiom, and in binary logic, it is a tautology. Yet the members of our imagined hominid band have taken up contradictory positions that must be resolved by providing reasons. They are compelled to engage in argumentation.

The reasons provided in the ensuing argument may not all be, at least prima facie, rationally related to the conclusion to take one course of action over another. Perhaps one side includes a feared warrior. In argumentation research, this is called the fallacy of the argumentum ad baculum: the argument from force.

SOCIAL REASONING: ARGUMENTATION

Although the study of argumentation spans many disciplines and theoretical perspectives, here we focus on the contributions of the new paradigm. This approach has primarily considered how and when so-called fallacies of informal argumentation are actually good arguments that should lead to changes in degrees of belief (Hahn & Oaksford 2007, Oaksford & Hahn 2004). The fallacies have accumulated in logic textbooks since Aristotle's *Sophistical Refutations*. They are arguments that are logically invalid, though persuasive.

By definition, the old-paradigm logical approach cannot account for these fallacies: They are fallacies because they are logically invalid. In MMT, if people find them persuasive they must be seen as victims of cognitive illusions (Khemlani & Johnson-Laird 2017). In contrast, the Bayesian new paradigm shows that, far from suffering an illusion, people often appropriately change their degree of belief in a conclusion given such arguments. This smooth generalization to the argumentative fallacies is, we suggest, one of the new paradigm's most important achievements, and one of the greatest challenges for the logical paradigm.

Bayesian Approaches to the Fallacies of Informal Argument

How can the fallacies sometimes be seen as justifiable? We consider four key examples: arguments from ignorance, ad hominem arguments, slippery slope arguments, and circular reasoning.

The argument from ignorance. Both propositions below are arguments from ignorance.

Proposition 2. Ghosts exist, because no one has proved that they do not.

Proposition 3. This drug is not toxic, because we can find no toxic effects.

Although Proposition 2 is intuitively weak, Proposition 3 is the standard mode of inference in product testing. The Bayesian approach distinguishes Propositions 2 and 3 by considering the degree of belief in the conclusion (C) to which an argument (a) should lead, using Bayes' theorem:

$$\Pr(\neg C|\neg a) = \frac{\Pr(\neg a|\neg C)\Pr(\neg C)}{\Pr(\neg a|\neg C)\Pr(\neg C) + \Pr(\neg a|C)\Pr(C)}. \qquad 1.$$

The strength of the conclusion in Proposition 3, $\Pr(\neg C|\neg a)$ (the drug is not toxic given no toxic effects were found), depends on the priors, $\Pr(C)$, and on the likelihoods, $\Pr(a|C)$ (i.e., sensitivity) and $\Pr(\neg a|\neg C)$ (i.e., specificity). The force of an argument, that is, how much our belief in the conclusion changes, can be indexed by the likelihood ratio $\Pr(a|C)/\Pr(a|\neg C)$. Proposition 2 is weaker than Proposition 3 because tests for toxicity are presumably far more reliable than tests for the nonexistence of ghosts. So, for toxicity, we obtain $\Pr(a|C) \gg \Pr(a|\neg C)$, and argument force is high; conversely, it seems hard to see how we could prove the nonexistence of ghosts, irrespective of whether they exist, so $\Pr(a|C) \approx \Pr(a|\neg C)$, and argument force is low.

Hahn & Oaksford (2007) argue that in typical real-world settings $\Pr(C|a) > \Pr(\neg C|\neg a)$, and consequently, positive arguments should be perceived as stronger than negative arguments. This prediction, and people's sensitivity to priors and likelihoods, has been confirmed experimentally (Hahn & Oaksford 2007, Hahn et al. 2005, Oaksford & Hahn 2004) and replicated across cultures (Karaaslan et al. 2018).

These experiments used a third-person argument evaluation paradigm. In this paradigm, participants are shown an argument between two interlocutors, say, Mary and John, and are told that Mary has a certain prior degree of belief in conclusion C (e.g., fairly convinced). John presents Mary with an argument, and participants must judge what degree of belief Mary should now have in C. The arguments are therefore placed in a social context where participants evaluate others' reactions to an argument.

The Bayesian model also captures the so-called epistemic closure version of the argument from ignorance: for example, that we may conclude that the book is in the library (C) because the online catalog does not say it is on loan, $\Pr(C|\neg\text{“}\neg C\text{”})$. This form is perhaps closer to Proposition 2. Modeling this case requires introducing a third possibility that the catalog says nothing (n) about whether the book is in the library or on loan. Most databases are not epistemically closed, in which case the probability of saying nothing is greater than zero: $\Pr(n|C) > 0$. When this is the case, not saying that the book is not in the library ($\neg\text{“}\neg C\text{”}$) is not equivalent to saying it is in the library (“C”), as the database may say nothing about the book's status. Experimentally varying the degree of closure [$\Pr(n|C)$] leads to results predicted by the Bayesian model (Hahn et al. 2005), with good quantitative data fits (Hahn & Oaksford 2007). People also regard Proposition 2 as weaker than the library example (Hahn & Oaksford 2007, experiment 3). Saying nothing (n) also arises in the "damned by faint praise" argument (Harris et al. 2013).

An important feature of argumentation is audience evaluation of the reliability of the person, newspaper, or government proposing the argument (Hahn et al. 2009, 2013). The probability that a source is reliable, $\Pr(R)$, can be factored into the likelihoods in Bayes' theorem [e.g., $\Pr(a|C)$ becomes $\Pr(a|C, R)\Pr(R) + \Pr(a|C, \neg R)(1 - \Pr(R))$, which caps argument strength (Bovens & Hartmann 2003). Manipulating source reliability has predictable effects on the strength of arguments from ignorance; higher reliability leads to greater argument force and hence stronger conclusion endorsement (Hahn & Oaksford 2007, Karaaslan et al. 2018).

The argument ad hominem. The ad hominem fallacy attacks the person proposing the argument, not the argument itself. In court, for example, the testimony of character witnesses is clearly viewed as having probative value. Moreover, the fallacy has more and less compelling versions, as here.

Proposition 4. Candidate Jones has no right to moralize about the family since

a. he was once seen arguing with his wife;
b. he cheats on his wife.

The continuation in part *b* of Proposition 4 is stronger than in part *a*. This fallacy has been modeled in a variety of ways (Bhatia & Oaksford 2015, Harris et al. 2012, Oaksford & Hahn 2013). We might initially view Jones's moralizing as credible if we assume by default that people are reliable. This assumption is undermined strongly by part *b* of Proposition 4, but barely at all by part *a*. Part *b* of Proposition 4 makes us question Jones's reliability and reduces the credibility of the conclusions of his moralizing (Oaksford & Hahn 2013). Experimentally varying reliability in different ways has shown the predicted effects on this argument (Bhatia & Oaksford 2015, Harris et al. 2012, Oaksford & Hahn 2013), and a Bayesian model provides good parameter-free fits to the data (Harris et al. 2012).

The slippery slope argument. The slippery slope argument (SSA), widely deployed from law to bioethics, also has weak and strong versions.

Proposition 5. If we allow gay marriage, then people will want to marry their pets.

Proposition 6. If voluntary euthanasia is legalized, then involuntary euthanasia will be.

Both have been advocated in popular debate. Although Proposition 5 is intuitively absurd, variants of Proposition 6 are influential. The Bayesian approach treats SSAs as consequentialist arguments based on a cost-benefit analysis (Corner et al. 2011, Hahn & Oaksford 2007): Do not act (p = allow gay marriage) unless the benefit associated with the action [$U(p)$] exceeds the expected cost [$\Pr(q|p)U(q)$] of the consequences to which it may lead (q = allowing interspecies marriage). In this respect, SSAs are similar to warnings or utility conditionals more generally (Bonnefon 2009) (see section titled Utility Conditionals). In contrast to Proposition 6, Proposition 5 is weak because whatever we think of the merits of interspecies marriage [$U(q)$], the probability that allowing gay marriage will have this consequence [$\Pr(q|p)$] is so low that the expected cost is negligible.

A crucial feature of SSAs is that the slipperiness of the slope [captured by $\Pr(q|p)$] may depend on the flexibility of our categories (Hahn & Oaksford 2007). Proposition 5 can be characterized as the argument that if we categorize gay marriage (a) as allowable (Fa), this will lead to us categorizing interspecies marriage (*b*) as allowable (*Fb*). Thus, the similarity between *a* and *b* should (and does) affect people's willingness to assign them to the same category and hence to endorse the corresponding SSA (Corner et al. 2011; see also Rai & Holyoak 2014 on moral hypocrisy).

Circular reasoning. Circularity (Hahn 2011) is an oddity among the fallacies because certain formulations are logically valid (*if p then p* is a tautology). As usual, the fallacy has strong and weak variants.

Proposition 7. God exists because the Bible says so and the Bible is the word of God.

Proposition 8. Electrons exist because they make 3-cm tracks in a cloud chamber.

Proposition 8 is a standard form of inductive inference in science concerning unobservables and is intuitively stronger than Proposition 7. A good argument should change our belief in the conclusion (Harman 1986). With vicious circularity (God exists because God exists), which simply restates the premise in the conclusion, no such change occurs: The conclusion simply inherits the credibility of the premise. Propositions 7 and 8, in contrast, are self-dependent but not vicious, and they can potentially change our degree of belief in a conclusion (captured in a hierarchical Bayesian model; see Hahn 2011, Hahn & Oaksford 2007). The perceived strength of a circular argument varies predictably with the number of possible interpretations (Hahn & Oaksford 2007, experiments 1 and 2), which explains our intuitions about the strengths of Propositions 7 and 8. There are more credible interpretations of who wrote the Bible (Proposition 7 is weak) than there are of what caused a 3-cm track in a cloud chamber (Proposition 8 is strong).

Argument schemes and Bayesian networks. Bayesian networks (Pearl 1988, 2000) provide a useful framework for representing not just circularity but a wide variety of argument schemes proposed in the argumentation literature (e.g., Walton et al. 2008). For example, in the argument from authority or expert opinion, an argument scheme might consist of questions like, "Is this expert opinion consistent with other expert opinions?" "Are the opinions independent?" and so on. Bayesian networks provide a natural way of implementing different argument schemes (Hahn & Hornikx 2016, Hahn et al. 2013, Harris et al. 2016) and a normative account of how to combine expert judgements, levels of expertise, and trustworthiness that fits the experimental data (Harris et al. 2016). The behavior of social networks of Bayesian agents has also been explored in modeling appeal to popular opinion and related arguments (e.g., Cook & Lewandowsky 2016, Hahn et al. 2018), further emphasizing the social dimension of reasoning in the new paradigm.

Further Approaches to Argumentation

Other researchers have addressed human reasoning from a social, argumentative perspective, reinterpreting apparent biases (e.g., Mercier & Sperber 2011, 2017). For example, confirmation bias makes sense from an argumentative standpoint: Like any courtroom lawyer, we will of course adduce evidence for, not against, the position we advocate (Mercier & Sperber 2011). In argument evaluation, as in the experiments we reported earlier, people are persuaded by rational arguments and not merely by the most confident group member (Trouche et al. 2014). Moreover, the correctness of initially counterintuitive reasoning rapidly persuades groups when they can argue (Claidière et al. 2017). However, when producing arguments, we may be biased and cognitively lazy (Trouche et al. 2016). This strand of argumentation research eschews formal normative frameworks, arguing that people's reasoning makes evolutionary sense on a case-by-case basis (Mercier & Sperber 2017). Nonetheless, Mercier has argued that some norms of reasoning used in argumentation are universal (Mercier 2011, Mercier et al. 2015; see also Karaaslan et al. 2018), so that differential levels of adherence to the principle of noncontradiction by Eastern and Western cultures (Peng & Nisbett 1999) may be illusory. In summary, both the argument-based and the Bayesian approach to reasoning agree that argumentation and its social function are primary, and they disagree only on particulars (e.g., the role of normative theories and cognitive modularity).

INDIVIDUAL REASONING: DEDUCTION AND INDUCTION

We have seen how fallacies of argumentation provide a typology of argument forms that can be strong or weak dependent on their content. Logical arguments, though often treated separately since Aristotle, are no different: Valid arguments, like MP, or logical fallacies, like affirming the consequent (AC) (i.e., *if p then q*; *q*, therefore *p*), can also be strong or weak depending on content and context (see Eva & Hartmann 2018). Reasoning is a subtype of argument, whose purpose is evaluating not logical validity but change in belief (Harman 1986). We begin with deductive reasoning, highlighting recent work in the new paradigm and referring to older research only when required.

Deductive Reasoning

We look at both propositional (*and*, *or*, *not*, *if*...*then*) and quantified (*all*, *some*) deductive reasoning. For the former we focus on the conditional *if*...*then*, looking also at reasoning with causal and utility conditionals.

Propositional reasoning. Patterns of reasoning about the logical connectives—*and*, *or*, *not*, and, in particular, *if*...*then*—have been intensively studied experimentally and theoretically. One approach is probability logic, which assigns probabilities to propositions so that they respect the laws of probability (Coletti & Scozzafava 2002, Pfeifer & Kleiter 2009). Such an assignment is probabilistically coherent, the probabilistic analog of logical consistency. In the new paradigm, researchers have investigated a set of closely related probabilistic concepts, including probabilistic validity, Bayesian and Jeffrey conditionalization, and dynamic inference, which we now introduce.

Probabilistic coherence. We first show how coherence applies to the MP inference in Proposition 1.

Proposition 9.

	If p then q	$\Pr(q\|p) = a$	(A′)
	p	$\Pr(p) = b$	(B′)
Therefore	q	$\Pr(q) = [ab, ab + (1 - b)]$	(C′)

As we have seen, the Equation (Edgington 1995) identifies the probability of a conditional, Pr(*if p then q*), with the conditional probability, $\Pr(q|p)$. The conditional probability is determined from background knowledge by the so-called Ramsey test: Add the antecedent to your stock of beliefs (assume that *the key is turned*), adjust to accommodate this belief, and read off the probability of the conclusion (*the car starts*) from your adjusted stock of beliefs. Thus, the conditional premise A′, *if I turn the key*, *the car starts*, has an associated degree of belief (a) determined by the Ramsey test. The categorical premise B′, *the key is turned*, is also assigned a degree of belief (b) depending on how it was learned. If we are uncertain of A′ or B′ (a and b are <1), then by the law of total probability, we can logically deduce a coherence interval for the probability [$\Pr(q)$] of the conclusion (C′) that *the car starts*:

$$\Pr(q) = \Pr(q|p)\Pr(p) + \Pr(q|\neg p)\left[1 - \Pr(p)\right]. \qquad 2.$$

This formula has a parameter, $\Pr(q|\neg p)$, the probability that *the car starts without turning the key*, not fixed by the premises. If $\Pr(q|\neg p) = 0$, then we have the lower bound, ab. If $\Pr(q|\neg p) = 1$, then we have the upper bound, $ab + (1 - b)$. Researchers have developed intuitive methods for calculating coherence intervals (Politzer 2016).

As we observed earlier (see section titled The Probabilistic Paradigm), the Equation has been experimentally confirmed (e.g., Evans et al. 2003, Oberauer & Wilhelm 2003, Politzer et al. 2010). In the old binary logical paradigm, the conditional merely rules out counterexamples (p, $\neg q$), so that $\Pr(\textit{if } p \textit{ then } q) = 1 - \Pr(p, \neg q) = \Pr(\neg p) + \Pr(p, q)$; this is never observed experimentally. The Equation also survives philosophical triviality arguments (Lewis 1976); experimental participants reject the key step in such arguments (Douven & Verbrugge 2013).

For MP, when a and b are varied, people's degrees of belief in the conclusion reliably fall into the relevant coherence interval at above-chance levels (Cruz et al. 2015, Evans et al. 2015, Pfeifer & Kleiter 2009, Politzer & Baratgin 2016, Singmann et al. 2014). This work also looked at MT and the logical fallacies, AC and denying the antecedent (DA) (i.e., *if p then q*; $\neg p$, therefore $\neg q$), which also have associated coherence intervals. Thus, it can be rational to endorse these fallacies to some degree. Coherence was above chance for DA, but not reliably for MT or AC. Cruz et al. (2015) suggested that this may arise from excessive task complexity and tested eight simple one-premise arguments (e.g., p or q; therefore, *if* $\neg p$ *then* q), finding above-chance coherence in each.

p-validity. The probabilistic analog of logical validity is p-validity (Adams 1998). As logical validity relies on consistency, so p-validity relies on coherence (Coletti & Scozzafava 2002, Pfeifer & Kleiter 2009). An argument is p-valid if and only if the uncertainty of the conclusion [$1 - \Pr(\textit{conclusion})$] cannot coherently be greater than the sum of the uncertainties of the premises: $\sum_i^n [1 - \Pr(\textit{premises}_i)]$. In Proposition 9, if $a = 0.9$ and $b = 0.5$, the sum of the premise uncertainties is 0.6, and at the lower bound of the coherence interval, the uncertainty of the conclusion is 0.55, while at the upper bound, it is 0.05. MP is p-valid. Although DA has a well-defined coherence interval, it is not p-valid. To see why, one should note that if the premise probabilities are both 1, then the sum of their uncertainties is 0; however, the uncertainty of the conclusion ranges between 0 and 1, that is, higher than the sum of the premise uncertainties.

The concept of p-validity is theoretically useful but perhaps not psychologically significant: People are interested in the degree to which an argument should change their degree of belief in a conclusion and not in whether it is valid (Harman 1986). However, p-validity is useful in discriminating between theories. Many rules of inference that are logically valid, like strengthening the antecedent, are not p-valid.

In classical binary logic, the conditional *if p then q* is interpreted as true unless both p is true and q is false. This material conditional interpretation generates a number of so-called paradoxes: inferences that are logically valid but counterintuitive. So, from *the moon is not made of cheese*, one can infer that *if the moon is made of cheese*, *then the car will start*, because the material conditional is true if its antecedent is false. In the old logical paradigm, the paradoxes are valid but pragmatically odd (Johnson-Laird & Byrne 2002); however, the paradoxes are not p-valid, and people do not endorse them (Pfeifer & Kleiter 2011). An inference of recent interest is *or*-introduction (p, therefore *p or q*), which is logically valid and p-valid. However, in the most recent version of MMT (e.g., Hinterecker et al. 2016) *or*-introduction is deemed invalid. This interpretation leads to seemingly absurd consequences (Baratgin et al. 2015), and people do endorse *or*-introduction (Cruz et al. 2015).

Conditionalization. According to Equation 2, if people learn the key was definitely turned [$\Pr(p) = 1$], then the probability of the conclusion is the conditional probability a, by Bayesian conditionalization, which yields a point value, not an interval. Earlier work in the new paradigm used Bayesian conditionalization (Oaksford et al. 2000) and assumed that people recruited information about $\Pr(q|\neg p)$ (or an equivalent value) from background knowledge to model MT and the fallacies. This approach provided good fits to most of the experimental results available at the time (Oaksford & Chater 2007).

Bayesian conditionalization assumes people have complete certainty [$\Pr(p) = 1$] on learning that the key is turned. People's original assessment of this probability, from prior knowledge, is that key-turning events are rare. Learning that the key was turned consequently changes the probability from, say, $\Pr_0(p) = 0.01$ to $\Pr_1(p) = 1$. Such certainty, however, is unusual: Typically, the categorical premise is at least somewhat uncertain—perhaps the key was not turned after all (Stevenson & Over 2001). In this case, Equation 2 provides a way of updating our degrees of belief to create the new distribution below through Jeffrey conditionalization (Jeffrey 2004):

$$\Pr_1(q) = \Pr_0(q|p)\Pr_1(p) + \Pr_0(q|\neg p)[1 - \Pr_1(p)]. \qquad 3.$$

Thus, $\Pr_1(q)$ depends on prior knowledge of $\Pr_0(q|\neg p)$, which is not specified in the premises; hence, the change of belief depends on background knowledge even for MP [when $\Pr_0(p) < \Pr_1(p) < 1$]. When participants provide estimates of all the relevant probabilities, their estimates of conclusion probabilities conform well to Jeffrey's rule (Zhao & Osherson 2010).

Dynamic inference. When learning some premise information changes some of the other premise probabilities, an inference is dynamic (Adams 1998; Gilio & Over 2012; Oaksford & Chater 2007, 2013). When we learn that $\Pr_1(p) > \Pr_0(p)$, we focus on a particular key-turning event; this does not alter the general probability of turning car keys, which remains low (Baratgin & Politzer 2010). However, learning some premises may require an update of beliefs about others (Baratgin & Politzer 2010; Oaksford & Chater 2007, 2013). As we have observed (see section titled Problems for the Old Paradigm), the MT premise, *the car did not start*, suggests a counterexample to an implicit conditional: *The key was turned, but the car did not start for some reason*. Following classical logic, we should reject the conditional premise as false. Although rejecting this useful action-guiding generalization violates common sense, some update to our degree of belief in the conditional is required (Oaksford & Chater 2007). Indeed, when explicitly faced with counterexamples, people reduce their degree of belief in the conditional premise from their initial assessment (Wolf et al. 2012).

In Equation 3, however, the conditional probabilities associated with the conditional premise are assumed to be invariant or rigid between the old ($\Pr_0$) and the new ($\Pr_1$) distribution, that is, $\Pr_1(q|p) = \Pr_0(q|p)$. People respect invariance on some tasks but not others (Zhao & Osherson 2010). The assumption that MT and the fallacies violate invariance also allows better fits to empirical asymmetries between MP and MT and between AC and DA (Oaksford & Chater 2007, 2013). On this account, the premises of MT and the fallacies provide potentially falsifying instances, causing people to update to a lower conditional probability, $\Pr_1(q|p) < \Pr_0(q|p)$, before they draw these inferences.

Dynamic inference is a more radical form of nonmonotonicity that goes beyond rejecting strengthening the antecedent. Just as we argued that logic could not tell us how to resolve contradictions, Bayesian or Jeffrey conditionalization cannot tell us what the new value of $\Pr_1(q|p)$ should be—indeed, both assume invariance. How, then, should people update their beliefs when invariance does not hold? One approach is to convert dynamic inference problems into an equivalent static problem (Gilio & Over 2012), though this may not be viable for the MT inference (Adams 1998, Oaksford & Chater 2013). Another option is to use Bayesian model discrimination to infer the new conditional probability (Oaksford & Chater 2007, 2013), but this approach seems unlikely to generalize.

A more general approach is to minimize a distance metric, such as the Kullback-Leibler (KL) distance from information theory, between the new (Pr_1) and the old (Pr_0) probability distributions (Douven 2012, Eva & Hartmann 2018, Oaksford & Chater 2013). Therefore, on learning a new premise probability, *the key was turned* [$\mathrm{Pr}_1(p) > \mathrm{Pr}_0(p)$], the other probabilities in the new distribution are chosen to be minimally different from the old distribution. If the premise is assumed to be certain [$\mathrm{Pr}_1(p) = 1$], this approach is equivalent to Bayesian conditionalization. The approach generalizes to changes in the probability of the conditional premise [e.g., $\mathrm{Pr}_1(q|p) \neq \mathrm{Pr}_0(q|p)$], and hence to a new degree of belief in the conclusion (Douven 2012, Eva & Hartmann 2018, Hartmann & Rafiee Rad 2012). For example, in MT (learning of a nonstarting car), we may come to suspect the car is less reliable than we thought [$\mathrm{Pr}_1(q|p) < \mathrm{Pr}_0(q|p)$].

Implications for dual processes. Distance minimization has implications for the status of the logically valid inferences and consequently for dual-process or dual-source theories of reasoning, which postulate a slow binary logical System 2 alongside a fast belief-based probabilistic System 1 (Heit & Rotello 2010, Klauer et al. 2010, Rips 2001, Stanovich 2011). The main evidence for dual systems of reasoning comes from experiments showing that although most inferential behavior depends on our probabilistic beliefs, people prefer logically valid arguments (Evans et al. 1983, Singmann et al. 2016). For example, a model that includes a binary logical component in which MP and MT are valid significantly improves the fit (from 80% to 86% of the variance) over Bayesian alternatives, including KL distance minimization (Singmann et al. 2016). However, there are probabilistic interpretations of MP and MT (unlike the fallacies DA and AC) such that they always raise the probability of the conclusion (see Eva & Hartmann 2018, theorem 5). Consequently, MP and MT are not only logically valid but also probabilistically reliable, and sensitivity to this latter distinction may explain the slight improvement in the model fits without positing a separate logical, analytic System 2. Perhaps, in short, it is probabilities all the way up (Elqayam & Evans 2011, Oaksford & Chater 2012).

In summary, the new-paradigm probabilistic approach can be purely deductive, deriving a coherence interval for the conclusion by simply using the information in the premises. We have suggested that research is required to account for dynamic inference, whereby people recruit more information than is given explicitly in the premises to update their beliefs about the conclusion. Thus, the new paradigm will need to address the challenge that people may consider more than probabilities when updating their beliefs (e.g., including factors like explanatory goodness; see Douven & Mirabile 2018, Douven & Schupbach 2015b).

Causal conditionals and inferential relevance. Much research over the last 30 years (Byrne 1989, Cummins et al. 1991) has focused on causal conditional reasoning: The antecedent event (turning the key) causes the consequent event (the car starting). Just as models of argumentation use causal Bayes nets (CBNs) (Hahn & Hornikx 2016), so do models of dynamic causal inference (Hartmann & Rafiee Rad 2012). In a CBN, a conditional is represented as an edge in a directed acyclic graph, whose corresponding causal power (W) (Cheng 1997) can be computed from $\mathrm{Pr}(q|p)$ and $\mathrm{Pr}(q|\neg p)$:

$$W = \frac{\mathrm{Pr}(q|p) - \mathrm{Pr}(q|\neg p)}{1 - \mathrm{Pr}(q|\neg p)}. \quad 4.$$

In a CBN with two alternative causes in a common effect graph (cause 1 $\rightarrow$ effect $\leftarrow$ cause 2), the effects are combined using the noisy-OR function (Pearl 1988).

CBNs have been used to model two key findings: suppression effects (Fernbach & Erb 2013, Oaksford & Chater 2017) and explaining away (Ali et al. 2011, Hall et al. 2016). If people are provided with another conditional describing an alternative cause—e.g., *if hot-wired, the car starts*—that creates a common-effect structure, they endorse AC less. That is, the AC inference is suppressed: Learning that the car started may not mean that the key was turned, as the car could have been hot-wired. The precise quantity that best captures suppression is debated (Cummins 2014, Fernbach & Erb 2013, Oaksford & Chater 2017), but CBNs better explain these effects than previous probabilistic models (e.g., Oaksford & Chater 2007). In particular, they explain why many implicitly present alternative causes (Cummins et al. 1991) may yield weaker suppression effects than a single explicit alternative (Byrne 1989). In the latter case, people explicitly represent the alternative cause and simulate the effects of its presence (Oaksford & Chater 2017).

To understand explaining away, we again assume the common-effect structure for two conditionals. When *the car started* is set to 1, as in AC, a CBN assigns a high probability to *the key was turned*, which drops when *the car was hot-wired* is set to 1. Hot-wiring explains away turning the key as the cause of the car starting, a pattern that is puzzling for the logical paradigm (Ali et al. 2011, Hall et al. 2016).

The paradoxes of material implication in classical logic highlight the importance of the relevance of premises and conclusions (whether causal or not) (Douven et al. 2018). According to material implication, *if the moon is made of green cheese, the car starts* is true merely in virtue of the falsity of its antecedent. However, this seems bizarre because antecedent and consequent seem utterly unconnected. The numerator of Cheng's (1997) causal power (see Equation 4), called delta-P (ΔP), provides a probabilistic index of inferential relevance (Skovgaard-Olsen et al. 2017). The probability that the car starts is unaffected by whether or not the moon is made of cheese, so $\Delta P = 0$ because $\Pr(q|p) = \Pr(q|\neg p)$. Following the Equation, early indications were that judgements of the probability of the conditional depend only on $\Pr(q|p)$ (Over et al. 2007). However, recent work suggests that inferential relevance plays a moderating role: The Equation holds when $\Delta P > 0$, when p and q are positively inferentially relevant, but much less so when they are irrelevant ($\Delta P = 0$) or negatively relevant ($\Delta P < 0$) (Skovgaard-Olsen et al. 2017). There is debate over whether relevance effects are semantic, part of conventionalized meaning (Krzyżanowska et al. 2017, Vidal & Baratgin 2017), or pragmatic (inferred from discourse context; see Cruz et al. 2016). Recent evidence suggests that positive relevance is semantic rather than a presupposition or conversational implicature (Skovgaard-Olsen et al. 2019), and a recent theoretical model of some novel experiments on inferential relevance may be compatible with CBNs (Douven et al. 2018). The logical paradigm lacks the resources even to frame these questions. The complex empirical picture highlights the pressing need for theoretical integration in the new paradigm, perhaps through CBNs.

Utility conditionals. Like slippery slope arguments, conditionals can focus on value and action. Thus, consequential (Proposition 10) and persuasion (Proposition 11) conditionals attempt to alter commitment to the antecedent action by manipulating the utility of the consequent event (Bonnefon & Hilton 2004).

Proposition 10. If the CEO admits fraud, he will be sent to jail.

Proposition 11. If the Kyoto accord is ratified, greenhouse gas emissions will be reduced.

Conditional inducements form another distinct class of utility conditionals, including conditional promises (Proposition 12), threats (Proposition 13), tips (Proposition 14), and warnings (Proposition 15) (e.g., Evans et al. 2008, Ohm & Thompson 2004).

Proposition 12. If you wash up for me, then I will give you some money.

Proposition 13. If you do not pay up, then I will send the debt collectors around.

Proposition 14. If you revise t-tests, then you will do well in the statistics exam.

Proposition 15. If you park in a private car park, then your car will get clamped.

Factors that distinguish these conditionals include whether the antecedent concerns a third party (Propositions 10 and 11) or the hearer (Propositions 12–15), whether the consequent action is by the speaker (Propositions 12 and 13) or not (Propositions 14 and 15), and whether the speaker values the antecedent positively (Proposition 12) or negatively (Proposition 13). These pragmatic factors determine how using these conditionals in conversation affects other people's actions or beliefs.

Bonnefon (2009) developed a general theory of utility conditionals, using utility grids to represent these pragmatic factors and utilities. The theory includes a set of folk axioms about how people use this information to make decisions, generating testable predictions about the conclusions invited by a utility conditional. The utility grids create the possibility of generating previously unexplored utility conditionals and of testing a random sample of these revealed effects predicted by the folk axioms (Bonnefon 2012). The theory has also been modified to incorporate potential causal linkage between antecedent and consequent (Bonnefon & Sloman 2013).

Utility conditionals appear to persuade via a philosophically contentious move: deriving an "ought" from an "is." For example, the conclusions we are invited to draw from the factual Propositions 10 and 11 are the modal propositions below (Elqayam et al. 2015).

Proposition 16. The CEO should not admit fraud.

Proposition 17. The Kyoto accord should be ratified.

However, an apparent is-to-ought puzzle is illusory: The relevant utilities are recruited from our background knowledge. These inferences presuppose that we ought to act to achieve our valued goals: Thus, positive value is transferred from the consequent (reducing greenhouse gases is good) to the antecedent (ratifying the Kyoto accord is good).

The strength of these inferences depends on whether the conclusion has positive (Proposition 11) or negative (Proposition 10) value, the causal strength of the link between antecedent and consequent, and defeasibility, as tested using the suppression paradigm (Elqayam et al. 2015). CBNs can model the latter two factors, again emphasizing their possible theoretical importance in the new paradigm. Moreover, influence diagrams (Howard & Matheson 1981) can incorporate background utilities in CBNs (Oaksford & Chater 2017).

The same principles help explain reasoning with deontic conditionals (about how people should behave): for example, *if you drink beer, you must be over 18*. An enforcer of this rule has the goal of finding cheaters who are drinking underage (Manktelow & Over 1991, Oaksford & Chater 1994), and this is sometimes misleadingly viewed as a logical search for counterexamples.

Directives are a further form of utility conditional, which are also deontic: for example, *if you see an unusual blip, then (you should) launch the depth charges* (Hilton et al. 2005). The weapons officer's goal in uttering this conditional is that an inferior perform the consequent action should the antecedent conditions arise. Goal structure can alter people's understanding of directives: Should the inferior primarily avoid misses (blips but no depth charges launched) or false alarms (no blip but depth charges launched)? The data show that these factors affect the information examined to check that the inferior has obeyed the directive; the preferred conditional formulation (e.g., avoid misses, *if p then q*, or avoid false alarms, *q only if p*); and, when these different formulations are used, people's comprehension of the implicit goal (Hilton et al. 2005). A directive only succeeds if it coordinates the inferior's actions with the officer's intentions. As Hilton et al. (2005, p. 403) observe, "Rationality, here, is thus social and pragmatic, determined by the successful coordination of the speaker and the hearer to achieve shared organisational goals."

Quantified inference. These inferences involve the logical quantifiers *all*, *some*, *some-not*, and *none*, and the generalized quantifiers *most* and *few*. Syllogisms are two-premise arguments using these terms, as for example in the logically valid syllogism below.

Proposition 18.

	All beekeepers (B) are artists (A).
	Some chemists (C) are beekeepers.
Therefore	Some C are A.

The probability heuristics model (PHM) (Chater & Oaksford 1999) assigns a probabilistic semantics to quantified assertions: *all*, $\Pr(y|x) = 1$; *some*, $\Pr(y|x) > 0$; *some-not*, $\Pr(y|x) < 1$; *none*, $\Pr(y|x) = 0$; *few*, $0 < \Pr(y|x) < \Delta$; and *most*, $1 - \Delta < \Pr(y|x) < 1$, where Δ is small. There are 64 possible premise pairs, each defining one of four dependency graphs determined by the figure [the placement of the end (A, C) and middle (B) terms]. For example, Proposition 18 defines the graph *chemists* → *beekeepers* → *artists*. The premises determine some of the parameters of the dependency graph: Here, $\Pr(A|B) = 1$ and $\Pr(B|C) > 0$. We can then prove whether the premise probabilities constrain the possible conclusion probabilities, $\Pr(C|A)$ or $\Pr(A|C)$. In this case, it turns out that $\Pr(A|C) > 0$ must hold given the premise probabilities, and so a *some C are A* conclusion follows p-validly. A syllogism is p-invalid when $\Pr(C|A)$ or $\Pr(A|C)$ are unconstrained by the premises and can take any value between 0 and 1. This notion of probabilistic validity is similar to probabilistic coherence, but it differs from it because dependency graphs introduce additional independence assumptions.

PHM uses simple heuristics, the two most important of which are the *min*- and the *max*-heuristic. The *min*-heuristic is defined over an ordering in informativeness of quantified statements, such that *all* > *most* > *few* > *some* > *none* ≫ *some-not*. This order is justified by the probabilistic semantics and a rarity assumption (i.e., most predicates apply to a small subset of objects). The *min*-heuristic states that the quantifier in the conclusion should be the same as the least informative premise (in Proposition 18, *some*). The *max*-heuristic states that confidence in this conclusion depends on the expected informativeness of the most informative premise (in Proposition 18, *all*). Expected informativeness is ordered such that *all* > *most* > *few* > *some* > *some-not* ≈ *none*. The expectation is calculated over the conclusions of all the p-valid syllogisms with that type of quantifier as its most informative premise.

PHM provided as good a fit to a meta-analysis of experiments using logical quantifiers as old-paradigm mental logic approaches (Rips 1994), but using fewer parameters (six versus nine) (Chater & Oaksford 1999). Comparing PHM to the other accounts of syllogistic reasoning available at the time involved identifying key predictions where PHM and these accounts diverged. In most cases, the results supported PHM. Two experiments also showed that PHM smoothly generalized to syllogisms using the quantifiers *most* and *few* (Chater & Oaksford 1999), which have still not been captured by nonprobabilistic accounts. MMT and PHM account equally well for experiments linking memory span to syllogistic reasoning (Copeland & Radvansky 2004), but PHM shows better fits for syllogisms with more than two premises (Copeland 2006), and it captures the distinction between strong and weak possible conclusions (Evans et al. 1999). In PHM, the *min*-heuristic conclusion is the strong conclusion (Oaksford & Chater 2007), whereas, in MMT, the strong conclusion follows in the first model constructed (Evans et al. 1999). However, for multiple-model syllogisms, MMT has no explanation of which model is constructed first or why it should conform to the *min*-heuristic (Oaksford & Chater 2007). A recent model comparison (Khemlani & Johnson-Laird 2012) concluded that no current theory adequately accounts for syllogistic reasoning.

The probabilistic representation model (PRM) (Hattori 2016) builds on both PHM and MMT. PRM assumes that people first construct a probabilistic prototype model (PPM). There are eight possible joint probabilities [e.g., $\Pr(A, \neg B, C)$] given the three terms in a syllogism, A, B, and C. To calculate these probabilities, PRM uses two parameters and assumes that end terms are independent given the middle term. One parameter is the probability of the middle term, $\Pr(B)$, which is the same for the end terms unless one term is *all*, when a coverage parameter (c) determines the degree of overlap between B and A or C. The premises may also imply that some of the joint probabilities will be 0. PRM then randomly samples from the resulting probability distribution—for example, an artist who is a chemist but not a beekeeper—consistent with the premises in Proposition 18. PRM represents the sampled individuals in a sample mental model (SMM). A sequential application of the *min*-heuristic then determines the conclusion. The *min*-heuristic selects a conclusion type, which PRM tests against the SMM to see if the conclusion is possible. The *some chemists are artists* conclusion in Proposition 18 is possible, as long as at least one of the sampled individuals is both a chemist and an artist. If no such individual is available, PRM checks the next quantifier down the order and so on. If no conclusion is consistent with the SMM, PRM reports no valid conclusion. Finally, PRM assesses the degree of confidence in the conclusion using the *max*-heuristic introducing four further free parameters.

PRM provides marginally better data fits than PHM or a parameterized MMT. Moreover, when varying the number of items sampled, six or seven items provide the best fit, consistent with the capacity of working memory. The model includes important theoretical innovations applicable in other areas of reasoning and decision making. First, PRM assumes people construct a generative PPM, which depends on the premises and on some assumptions about the world these premises describe. PRM adopts a set size-balancing principle, assuming that the size of the categories defined by the terms A, B, and C is the same or symmetrical and therefore the coverage parameter is high (best-fit value of $c = 0.91$) (Hattori 2016). There is evidence that people make this assumption, which can often lead to errors, and that symmetry inferences are uniquely human (Hattori 2016, Hattori & Oaksford 2007). Second, people sample this generative model to create a discrete record of each sampled item (SMM) (see also Oaksford & Chater 2017, Oaksford & Hall 2016). This limited sample in memory forms the basis for predicting the conclusion, in a similar way to how sampling models in judgement and decision making can predict people's decisions (e.g., Vul et al. 2014).

Inductive Reasoning

The psychology of reasoning has traditionally emphasized deductive reasoning, where logic seems most naturally applicable. However, from the standpoint of the new paradigm, reasoning is typically probabilistic and knowledge rich. Thus, so-called inductive inferences—in which conclusions extrapolate, often dramatically, beyond the information given in the premises—are continuous with so-called deductive inference. Almost every aspect of perception and cognition can, in a broad sense, be viewed as involving induction, and the relevant literatures in psychology, philosophy, statistics, and AI are vast. Here we briefly review work that focuses on verbal inductive reasoning.

Property induction. Paradigmatic verbal inductive arguments involve generalizing a property from one or more cases to new cases. For example, we might wonder which of the following conclusions follows most strongly:

Proposition 19.

	Mice have property X.
Therefore	Rats/cats have property X.

Intuitively, similarity plays a role: Mice are more similar to rats than to cats and hence seem likely to share properties. Recent work finds (e.g., Kemp & Tenenbaum 2009), though, that the nature of the property, conjoined with background knowledge, is crucial. If property X is having a virus, and background knowledge tells us that cats may catch the virus by eating mice, then the generalization to cats may be stronger.

Researchers have also considered multiple premises:

Proposition 20.

	Mice have property X.
	Rats/fish have property X.
Therefore	Dogs have property X.

Here, the diversity of examples seems persuasive (e.g., Heit & Feeney 2005). If mice and rats have property X, we may suspect it is limited to rodents; if mice and fish do, we may suspect X is very widespread across animals and most likely includes dogs. Again, though, the specific property and background knowledge are crucial. As for argumentation and deduction, such inferences are well modeled by hierarchical Bayesian models, which capture such knowledge (Kemp & Tenenbaum 2009).

Abduction. Abductive reasoning seeks to infer an explanation for given premises, sometimes known as inference to the best explanation (Harman 1965). Scientists, adult participants, and children seem to prefer simple explanations that explain as much as possible (Lombrozo 2016), and studies with children and adults indicate that they prefer explanations that have the fewest assumptions and are best integrated (Lombrozo & Vasilyeva 2017). For example, participants were provided with explanations for weight loss and fatigue, either separately in terms of loss of appetite and insomnia or via the additional root cause of depression: The latter, more integrated, explanation was preferred.

Explanations with broader coverage are also preferred. As with category induction, explanations that capture diverse evidence are preferred (Kim & Keil 2003); indeed, participants prefer explanations when they have been prompted to consider the variety of evidence they can account for (Preston & Epley 2005). Moreover, children between three and five years of age also show enriched inferences about the internal structure of a simple machine when prompted to generate verbal explanations.

Simplicity and explanatory breadth can be naturally captured within a probabilistic framework if there is prior bias to simple explanations—indeed, preferences for simple versus probable explanations are identical under some assumptions (e.g., Chater 1996). Whether Bayesian explanation entirely underpins other explanatory virtues continues to be debated (Douven & Schupbach 2015a).

Data selection. Inquiry into the world involves actively choosing which data to consider. According to logic-based approaches in the psychology of reasoning (e.g., Wason 1968), people should seek falsifications of their conjectures. However, from a probabilistic standpoint, people should select data expected to provide the most (or most useful) information. These approaches can make strikingly different predictions. For example, when evaluating a conditional such as *if you eat tripe, you will be ill* (on the assumption the tripe eating and illness are both rare), it seems natural to check whether any known tripe eaters became ill and whether any ill people ate tripe (i.e., favoring confirming, not falsifying, instances of the rule). However, it seems futile to survey non-ill people to check whether they have eaten tripe: If they have, this will be an informative falsification, but it is overwhelmingly likely that they have not, and hence no useful information will be gained.

Bayesian optimal data selection formalizes these intuitions and provides a rational analysis (Anderson 1991) of what was previously viewed as confirmation bias in data-selection tasks, most famously Wason's (1968) card-selection task. Variations of the conditional, by modifying rarity (as noted above) or introducing negations in the antecedent and/or conclusion, yield different predictions according to the Bayesian framework, but not according to the logical viewpoint (Oaksford & Chater 1994, 2007). Moreover, different underlying probabilistic models can be explored (Klauer 1999, Oaksford & Chater 1994; for a recent MMT perspective, see Ragni et al. 2018), as well as sequential information search in both adults and children (e.g., Nelson et al. 2014; cf. Coenen et al. 2018). The appropriate measure of amount of information is an open question, on both axiomatic and empirical grounds (e.g., Crupi et al. 2018).

Deductive and inductive reasoning: one process or two? The new paradigm assimilates deductive and inductive reasoning as probabilistic, knowledge-rich, and context-dependent arguments. This viewpoint seems incompatible with dual-process accounts in which a slow, logic-based reasoning system runs alongside a fast, associative, probabilistic system (e.g., Heit & Rotello 2010, Klauer et al. 2010, Rips 2001, Stanovich 2011). Recent work using signal detection theory tests directly whether one or two factors explain inductive and deductive inference, concluding that the data favor a unified account (Stephens et al. 2018).

FUTURE DIRECTIONS

We first establish further connections between the new paradigm and other areas of the cognitive and brain sciences, as these are likely to feature prominently in its future development. We then look briefly at the picture of the cognitive systems this review suggests and outline key challenges for its future development.

Verbal Reasoning as Part of the Cognitive and Brain Sciences

The new paradigm in the psychology of reasoning has a number of links with accounts of rationality in psychology, philosophy, neuroscience, and AI. The emphasis on probability theory links with the explosion of interest for Bayesian models of perception, motor control, language processing, and knowledge representation in psychology (e.g., Chater & Oaksford 2008). In particular, there are strong connections to Bayesian brain theory (Friston 2009, Oaksford & Hall 2016). The emphasis on knowledge-rich reasoning embedded in social interactions links with inferential theories of communication (Sperber & Wilson 1986). Moreover, the focus on the division of cognitive labor fits with developments in social epistemology (e.g., Goldman 1999), accounts of the collective rationality of simple agents in finance theory and economics (Farmer et al. 2005), and theories of cumulative cultural evolution (e.g., Richerson & Boyd 2005; see Karaaslan et al. 2018 on cultural differences in gullibility in intergenerational knowledge transmission).

The field of verbal reasoning sprang from the project of understanding how far human reasoning aligns with logic. In the new paradigm, argumentation, deduction, and induction alike are modeled within a social, probabilistic framework, and the rationale for seeing verbal reasoning as distinct from the rest of cognition has diminished. What makes verbal reasoning special may not be its underpinning cognitive machinery, but rather the fact that it is public, and hence open for debate and challenge (Mercier & Sperber 2017, Oaksford & Hall 2016).

Cognitive Theory and Its Challenges

A social perspective on reasoning may alleviate the need for individual, globally consistent knowledge bases in long-term memory (LTM) to support reasoning (Oaksford & Chater 2012), which is consistent with the fact that human knowledge is full of both gaps (e.g., Rozenblit & Keil 2002) and contradictions (Oaksford & Chater 1991). Consistency and coherence may be primarily at issue in particular adversarial exchanges when we make public commitments on which we are challenged. Individual reasoning may only involve representing locally coherent, small-scale generative models of the premises, like CBNs or PPMs. The process of constructing these models probably involves highly flexible, analogical generalizations from concrete exemplars in LTM rather than general principles (e.g., Chater 2018). Estimating probabilities from these models involves sampling the model (Vul et al. 2014) and occasionally creating discrete representations of the samples (Hattori 2016, Oaksford & Chater 2017). Most of this review has treated the Bayesian framework at the computational level. However, it is possible to view the proposals concerning CBNs, implicated in many areas of verbal reasoning, at the algorithmic level, as specifying the mental processes involved in a passage of reasoning. In argumentation, where people disagree, each interlocutor is attempting to align their discourse partner's local model with their own. Doing so requires an understanding of what is common knowledge between interlocutors; an argument cannot begin without agreement on interpretation that requires such knowledge.

Theories of reasoning must, therefore, confront some significant challenges. For example, in AI, the processes of reasoning over very large knowledge bases of fragmentary and inconsistent beliefs, which nonetheless can support the construction of locally coherent models, are not well understood (Pearl 1988). However, as we have suggested, such coherence is probably achieved by analogy to specific exemplars rather than by maintaining globally consistent individual knowledge bases. Moreover, the pragmatic processes involved in the interpretation of the premises and the conclusion may often involve far more sophisticated reasoning than the verbal inference under study (Sperber & Wilson 1986). Common knowledge and related notions of common belief are also extremely subtle (Clark 1996). Indeed, game-theoretic treatments of both common knowledge and communication are highly sophisticated and computationally complex (e.g., Aumann 1976).

CONCLUSION

Verbal reasoning is a social activity in which people communicate and argue. Individual knowledge is partial and fragmentary; reasoning must deal with probability and the integration of knowledge across people, via communication and argumentation. Traditional models of reasoning ignore uncertainty and communicative context and use formal logic as a standard for good reasoning. The new paradigm in the psychology of reasoning, by contrast, attempts to explain argumentation, deduction, and induction in a probabilistic framework, in continuity with models that have become widespread in the cognitive and brain sciences. Moreover, if collective, rather than individual, intelligence is primary (Sloman & Fernbach 2017), a key aspect of verbal reasoning may be in supporting public argument and debate (Mercier & Sperber 2017). While empirical and theoretical challenges remain, the new paradigm in the psychology of reasoning has yielded considerable progress and provides an integrative framework for future research.

FUTURE ISSUES

1. How can we embed reasoning in social-communicative contexts?
2. How can we model pragmatic reasoning in communicative exchanges?
3. Are causal and noncausal dependencies (e.g., reasons) modeled in the same way?
4. When do networks of individual reasoners synthesize knowledge rather than herding?
5. How does the fragmentary nature of individual knowledge support argumentation?
6. When do people reason, search for new information, or rely on other people (or Google)?
7. Is human verbal reasoning continuous with inferences underpinning perception and action?
8. Is there a Bayesian account of explanatory virtues, like integration, minimal assumptions, and goodness?
9. What are the origins of, and how do we reason with, common knowledge?

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

N.C. was supported by the ESRC Network for Integrated Behavioural Science (grant number ES/P008976/1).

LITERATURE CITED

Adams EW. 1998. *A Primer of Probability Logic*. Stanford, CA: CSLI Publ.

Ali N, Chater N, Oaksford M. 2011. The mental representation of causal conditional inference: causal models or mental models. *Cognition* 119:403–18

Anderson JR. 1991. Is human cognition adaptive? *Behav. Brain Sci.* 14:471–517

Aumann RJ. 1976. Agreeing to disagree. *Ann. Stat.* 4:1236–39

Baratgin J, Douven I, Evans JStBT, Oaksford M, et al. 2015. The new paradigm and mental models. *Trends Cogn. Sci.* 19:547–48

Baratgin J, Politzer G. 2010. Updating: a psychologically basic situation of probability revision. *Think. Reason.* 16:253–87

Bennett J. 1989 (1964). *Rationality*. Indianapolis, IN: Hackett Publ.

Bhatia J-S, Oaksford M. 2015. Discounting testimony with the argument ad hominem and a Bayesian congruent prior model. *J. Exp. Psychol. Learn. Mem. Cogn.* 41:1548–59

Bonnefon J-F. 2009. A theory of utility conditionals: paralogical reasoning from decision-theoretic leakage. *Psychol. Rev.* 116:888–907

Bonnefon J-F. 2012. Utility conditionals as consequential arguments: a random sampling experiment. *Think. Reason.* 18:379–93

Bonnefon J-F, Hilton DJ. 2004. Consequential conditionals: invited and suppressed inferences from valued outcomes. *J. Exp. Psychol. Learn. Mem. Cogn.* 30:28–37

Bonnefon J-F, Sloman SA. 2013. The causal structure of utility conditionals. *Cogn. Sci.* 37:193–209

Bovens L, Hartmann S. 2003. *Bayesian Epistemology*. Oxford, UK: Oxford Univ. Press

Byrne RM. 1989. Suppressing valid inferences with conditionals. *Cognition* 31:61–83
Carnap R. 1945. On inductive logic. *Philos. Sci.* 12:72–97
Chater N. 1996. Reconciling simplicity and likelihood principles in perceptual organization. *Psychol. Rev.* 103:566–81
Chater N. 2018. *The Mind Is Flat*. London: Allen Lane
Chater N, Oaksford M. 1999. The probability heuristics model of syllogistic reasoning. *Cogn. Psychol.* 38:191–258
Chater N, Oaksford M, eds. 2008. *The Probabilistic Mind: Prospects for Bayesian Cognitive Science*. Oxford, UK: Oxford Univ. Press
Cheng PW. 1997. From covariation to causation: a causal power theory. *Psychol. Rev.* 104:367–405
Claidière N, Trouche E, Mercier H. 2017. Argumentation and the diffusion of counter-intuitive beliefs. *J. Exp. Psychol. Gen.* 146:1052–66
Clark HH. 1996. *Using Language*. Cambridge, UK: Cambridge Univ. Press
Coenen A, Nelson JD, Gureckis TM. 2018. Asking the right questions about the psychology of human inquiry: nine open challenges. *Psychon. Bull. Rev.* In press. **https://doi.org/10.3758/s13423-018-1470-5**
Coletti G, Scozzafava R. 2002. *Probabilistic Logic in a Coherent Setting*. Dordrecht, Neth.: Kluwer
Cook J, Lewandowsky S. 2016. Rational irrationality: modeling climate change belief polarization using Bayesian networks. *Top. Cogn. Sci.* 8:160–79
Copeland DE. 2006. Theories of categorical reasoning and extended syllogisms. *Think. Reason.* 12:379–412
Copeland DE, Radvansky GA. 2004. Working memory and syllogistic reasoning. *Q. J. Exp. Psychol. Hum. Exp. Psychol.* 57A:1437–57
Corner A, Hahn U, Oaksford M. 2011. The psychological mechanism of the slippery slope argument. *J. Mem. Lang.* 64:133–52
Crupi V, Nelson JD, Meder B, Cevolani G, Tentori K. 2018. Generalized information theory meets human cognition: introducing a unified framework to model uncertainty and information search. *Cogn. Sci.* 42:1410–56
Cruz N, Baratgin J, Oaksford M, Over DE. 2015. Reasoning with ifs and ands and ors. *Front. Psychol.* 6:192
Cruz N, Over DE, Oaksford M, Baratgin J. 2016. Centering and the meaning of conditionals. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, ed. A Papafragou, D Grodner, D Mirman, JC Trueswell, pp. 1104–9. Austin, TX: Cogn. Sci. Soc.
Cummins DD. 2014. The impact of disablers on predictive inference. *J. Exp. Psychol. Learn. Mem. Cogn.* 40:1638–55
Cummins DD, Lubart T, Alksnis O, Rist R. 1991. Conditional reasoning and causation. *Mem. Cogn.* 19:274–82
Douven I. 2012. Learning conditional information. *Mind Lang.* 27:239–63
Douven I, Elqayam S, Singmann H, van Wijnbergen-Huitink J. 2018. Conditionals and inferential connections: a hypothetical inferential theory. *Cogn. Psychol.* 101:50–81
Douven I, Mirabile P. 2018. Best, second-best, and good-enough explanations: how they matter to reasoning. *J. Exp. Psychol. Learn. Mem. Cogn.* 44:1792–813
Douven I, Schupbach JN. 2015a. Probabilistic alternatives to Bayesianism: the case of explanationism. *Front. Psychol.* 6:459
Douven I, Schupbach JN. 2015b. The role of explanatory considerations in updating. *Cognition* 142:299–311
Douven I, Verbrugge S. 2013. The probabilities of conditionals revisited. *Cogn. Sci.* 37:711–30
Edgington D. 1995. On conditionals. *Mind* 104:235–329
Elqayam S, Evans JStBT. 2011. Subtracting "ought" from "is": descriptivism versus normativism in the study of human thinking. *Behav. Brain Sci.* 34:233–48
Elqayam S, Thompson VA, Wilkinson MR, Evans JStBT, Over DE. 2015. Deontic introduction: a theory of inference from is to ought. *J. Exp. Psychol. Learn. Mem. Cogn.* 41:1516–32
Eva B, Hartmann S. 2018. Bayesian argumentation and the value of logical validity. *Psychol. Rev.* 125(5):806–21
Evans JStBT. 2017. Belief bias in deductive reasoning. In *Cognitive Illusions: Intriguing Phenomena in Thinking, Judgment and Memory*, ed. RF Pohl, pp. 165–81. New York: Routledge. 2nd ed.

Evans JStBT, Barston JL, Pollard P. 1983. On the conflict between logic and belief in syllogistic reasoning. *Mem. Cogn.* 11:295–306

Evans JStBT, Handley SJ, Harper CNJ, Johnson-Laird PN. 1999. Reasoning about necessity and possibility: a test of the mental model theory of deduction. *J. Exp. Psychol. Learn. Mem. Cogn.* 25:1495–513

Evans JStBT, Handley SJ, Over DE. 2003. Conditionals and conditional probability. *J. Exp. Psychol. Learn. Mem. Cogn.* 29:321–35

Evans JStBT, Neilens H, Handley SJ, Over DE. 2008. When can we say "if"? *Cognition* 108:100–16

Evans JStBT, Thompson V, Over DE. 2015. Uncertain deduction and conditional reasoning. *Front. Psychol.* 6:398

Farmer JD, Patelli P, Zovko II. 2005. The predictive power of zero intelligence in financial markets. *PNAS* 102:2254–59

Fernbach PM, Erb CD. 2013. A quantitative causal model theory of conditional reasoning. *J. Exp. Psychol. Learn. Mem. Cogn.* 39:1327–43

Fodor JA. 1975. *The Language of Thought*. New York: Thomas Crowell

Friston KJ. 2009. The free-energy principle: a rough guide to the brain? *Trends Cogn. Sci.* 13:293–301

George C. 1995. The endorsement of the premises: assumption-based or belief-based reasoning. *Br. J. Psychol.* 86:93–111

Gilio A, Over D. 2012. The psychology of inferring conditionals from disjunctions: a probabilistic study. *J. Math. Psychol.* 56:118–31

Goldman AI. 1999. *Knowledge in a Social World*. Oxford, UK: Oxford Univ. Press

Hahn U. 2011. The problem of circularity in evidence, argument, and explanation. *Perspect. Psychol. Sci.* 6:172–82

Hahn U, Hansen JU, Olsson EJ. 2018. Truth tracking performance of social networks: how connectivity and clustering can make groups less competent. *Synthese*. In press. **https://doi.org/10.1007/s11229-018-01936-6**

Hahn U, Harris AJL, Corner AJ. 2009. Argument content and argument source: an exploration. *Informal Log.* 29:337–67

Hahn U, Hornikx J. 2016. A normative framework for argument quality: argumentation schemes with a Bayesian foundation. *Synthese* 193:1833–73

Hahn U, Oaksford M. 2007. The rationality of informal argumentation: a Bayesian approach to reasoning fallacies. *Psychol. Rev.* 114:704–32

Hahn U, Oaksford M, Bayindir H. 2005. How convinced should we be by negative evidence? In *Proceedings of the 27th Annual Conference of the Cognitive Science Society*, ed. B Bara, L Barsalou, M Bucciarelli, pp. 887–92. Mahwah, NJ: Erlbaum

Hahn U, Oaksford M, Harris AJL. 2013. Testimony and argument: a Bayesian perspective. In *Bayesian Argumentation*, ed. F Zenker, pp. 15–38. Dordrecht, Neth.: Springer

Hall S, Ali N, Chater N, Oaksford M. 2016. Discounting and augmentation in causal conditional reasoning: causal models or shallow encoding. *PLOS ONE* 11(12):e0167741

Harman G. 1965. The inference to the best explanation. *Philos. Rev.* 64:88–95

Harman G. 1986. *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press

Harris AJL, Corner A, Hahn U. 2013. James is polite and punctual (and useless): a Bayesian formalisation of faint praise. *Think. Reason.* 19:414–29

Harris AJL, Hahn U, Madsen JK, Hsu AS. 2016. The appeal to expert opinion: quantitative support for a Bayesian network approach. *Cogn. Sci.* 40:1496–533

Harris AJL, Hsu AS, Madsen JK. 2012. Because Hitler did it! Quantitative tests of Bayesian argumentation using *ad hominem*. *Think. Reason.* 18:311–43

Hartmann S, Rafiee Rad S. 2012. *Updating on conditionals = Kullback-Leibler distance + causal structure*. Paper presented at the Biennial Meeting of the Philosophy of Science Association, San Diego, Nov. 15–17

Hattori M. 2016. Probabilistic representation in syllogistic reasoning: a theory to integrate mental models and heuristics. *Cognition* 157:296–320

Hattori M, Oaksford M. 2007. Adaptive non-interventional heuristics for covariation detection in causal induction: model comparison and rational analysis. *Cogn. Sci.* 31:765–814

Heit E, Feeney A. 2005. Relations between premise similarity and inductive strength. *Psychon. Bull. Rev.* 12:340–44

Heit E, Rotello CM. 2010. Relations between inductive reasoning and deductive reasoning. *J. Exp. Psychol. Learn. Mem. Cogn.* 36:805–12

Hilton DJ. 1995. The social context of reasoning: conversational inference and rational judgment. *Psychol. Bull.* 118:248–71

Hilton DJ, Kemmelmeier M, Bonnefon J-F. 2005. Putting ifs to work: goal-based relevance in conditional directives. *J. Exp. Psychol. Gen.* 134:388–405

Hinterecker T, Knauff M, Johnson-Laird PN. 2016. Modality, probability, and mental models. *J. Exp. Psychol. Learn. Mem. Cogn.* 42:1606–20

Howard RA, Matheson JE. 1981. Influence diagrams. In *Readings on the Principles and Applications of Decision Analysis*, Vol. 2, ed. RA Howard, JE Matheson, pp. 719–62. Menlo Park, CA: Strateg. Decis. Group

Howson C, Urbach P. 1989. *Scientific Reasoning: The Bayesian Approach*. LaSalle, IL: Open Court

Jeffrey RC. 2004. *Subjective Probability: The Real Thing*. Cambridge, UK: Cambridge Univ. Press

Johnson-Laird PN. 1983. *Mental Models*. Cambridge, UK: Cambridge Univ. Press

Johnson-Laird PN, Byrne RMJ. 2002. Conditionals: a theory of meaning, pragmatics, and inference. *Psychol. Rev.* 109:646–78

Karaaslan H, Hohenberger A, Demir H, Hall S, Oaksford M. 2018. Cross-cultural differences in informal argumentation: norms, inductive biases and evidentiality. *J. Cogn. Cult.* 18:358–89

Kemp C, Tenenbaum JB. 2009. Structured statistical models of inductive reasoning. *Psychol. Rev.* 116:20–58

Khemlani SS, Johnson-Laird PN. 2012. Theories of the syllogism: a meta-analysis. *Psychol. Bull.* 138:427–57

Khemlani SS, Johnson-Laird PN. 2017. Illusions in reasoning. *Minds Mach*. 27:11–35

Kim NS, Keil FC. 2003. From symptoms to causes: diversity effects in diagnostic reasoning. *Mem. Cogn.* 31:155–65

Klauer KC. 1999. On the normative justification for information gain in Wason's selection task. *Psychol. Rev.* 106:215–22

Klauer K, Beller S, Hutter M. 2010. Conditional reasoning in context: a dual-source model of probabilistic inference. *J. Exp. Psychol. Learn. Mem. Cogn.* 36:298–323

Krzyżanowska KH, Collins PJ, Hahn U. 2017. Between a conditional's antecedent and its consequent: discourse coherence versus probabilistic relevance. *Cognition* 164:199–205

Lewis DK. 1976. Probabilities of conditionals and conditional probabilities. *Philos. Rev.* 85:297–315

Lombrozo T. 2016. Explanatory preferences shape learning and inference. *Trends Cogn. Sci.* 20:748–59

Lombrozo T, Vasilyeva N. 2017. Causal explanation. In *Oxford Handbook of Causal Reasoning*, ed. M Waldmann, pp. 415–32. Oxford, UK: Oxford Univ. Press

Manktelow KI, Over DE. 1991. Social roles and utilities in reasoning with deontic conditionals. *Cognition* 39:85–105

Mercier H. 2011. On the universality of argumentative reasoning. *J. Cogn. Cult.* 11:85–113

Mercier H, Sperber D. 2011. Why do humans reason? Arguments for an argumentative theory. *Behav. Brain Sci.* 34:57–74

Mercier H, Sperber D. 2017. *The Enigma of Reason*. Cambridge, MA: Harvard Univ. Press

Mercier H, Zhang J, Qu Y, Lu P, Van der Henst J. 2015. Do Easterners and Westerners treat contradiction differently? *J. Cogn. Cult.* 15:45–63

Millikan RG. 2006. Styles of rationality. In *Rational Animals?*, ed. S Hurley, M Nudds, pp. 117–26. Oxford, UK: Oxford Univ. Press

Nelson JD, Divjak B, Gudmundsdottir G, Martignon L, Meder B. 2014. Children's sequential information search is sensitive to environmental probabilities. *Cognition* 130:74–80

Oaksford M, Chater N. 1991. Against logicist cognitive science. *Mind Lang*. 6:1–38

Oaksford M, Chater N. 1994. A rational analysis of the selection task as optimal data selection. *Psychol. Rev.* 101:608–31

Oaksford M, Chater N. 2007. *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford, UK: Oxford Univ. Press

Oaksford M, Chater N. 2012. Dual processes, probabilities, and cognitive architecture. *Mind Soc*. 11:15–26

Oaksford M, Chater N. 2013. Dynamic inference and everyday conditional reasoning in the new paradigm. *Think. Reason.* 19:346–79

Oaksford M, Chater N. 2017. Causal models and conditional reasoning. In *Oxford Handbook of Causal Reasoning*, ed. M Waldmann, pp. 327–46. Oxford, UK: Oxford Univ. Press

Oaksford M, Chater N, Larkin J. 2000. Probabilities and polarity biases in conditional inference. *J. Exp. Psychol. Learn. Mem. Cogn.* 26:883–99

Oaksford M, Hahn U. 2004. A Bayesian analysis of the argument from ignorance. *Can. J. Exp. Psychol.* 58:75–85

Oaksford M, Hahn U. 2013. Why are we convinced by the ad hominem argument? Bayesian source reliability and pragma-dialectical discussion rules. In *Bayesian Argumentation: The Practical Side of Probability*, ed. F Zenker, pp. 39–58. New York: Springer Sci.

Oaksford M, Hall S. 2016. On the source of human irrationality. *Trends Cogn. Sci.* 20:336–44

Oberauer K, Wilhelm O. 2003. The meaning(s) of conditionals: conditional probabilities, mental models, and personal utilities. *J. Exp. Psychol. Learn. Mem. Cogn.* 29:680–93

Ohm E, Thompson VA. 2004. Everyday reasoning with inducements and advice. *Think. Reason.* 10:241–72

Over DE. 2009. New paradigm psychology of reasoning. *Think. Reason.* 15:431–38

Over DE, Hadjichristidis C, Evans JStBT, Handley SJ, Sloman SA. 2007. The probability of causal conditionals. *Cogn. Psychol.* 54:62–97

Pearl J. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufman

Pearl J. 2000. *Causality*. Cambridge, UK: Cambridge Univ. Press

Peng K, Nisbett RE. 1999. Culture, dialectics, and reasoning about contradiction. *Am. Psychol.* 54:741–54

Pfeifer N, Kleiter GD. 2009. Framing human inference by coherence based probability logic. *J. Appl. Log.* 7:206–17

Pfeifer N, Kleiter GD. 2011. Uncertain deductive reasoning. In *The Science of Reason: A Festschrift for Jonathan St. B. T. Evans*, ed. K Manktelow, DE Over, S Elqayam, pp. 145–66. Hove, UK: Psychol. Press

Politzer G. 2016. Deductive reasoning under uncertainty: a water tank analogy. *Erkenntnis* 81:479–506

Politzer G, Baratgin J. 2016. Deductive schemas with uncertain premises using qualitative probability expressions. *Think. Reason.* 22:78–98

Politzer G, Carles L. 2001. Belief revision and uncertain reasoning. *Think. Reason.* 7:217–34

Politzer G, Over DE, Baratgin J. 2010. Betting on conditionals. *Think. Reason.* 16:172–97

Preston J, Epley N. 2005 Explanations versus applications: the explanatory power of valuable beliefs. *Psychol. Sci.* 16:826–32

Ragni M, Kola I, Johnson-Laird PN. 2018. On selecting evidence to test hypotheses: a theory of selection tasks. *Psychol. Bull.* 144:779–96

Rai TS, Holyoak KJ. 2014. Rational hypocrisy: a Bayesian analysis based on informal argumentation and slippery slopes. *Cogn. Sci.* 38:1456–67

Richerson PJ, Boyd R. 2005. *Not by Genes Alone: How Culture Transformed Human Evolution*. Chicago: Univ. Chicago Press

Rips LJ. 1994. *The Psychology of Proof: Deductive Reasoning in Human Thinking*. Cambridge, MA: MIT Press

Rips LJ. 2001. Two kinds of reasoning. *Psychol. Sci.* 12:129–34

Rozenblit L, Keil F. 2002. The misunderstood limits of folk science: an illusion of explanatory depth. *Cogn. Sci.* 26:521–62

Salmon WC. 1967. Carnap's inductive logic. *J. Philos.* 64:725–39

Singmann H, Klauer KC, Beller S. 2016. Probabilistic conditional reasoning: disentangling form and content with the dual-source model. *Cogn. Psychol.* 88:61–87

Singmann H, Klauer KC, Over DE. 2014. New normative standards of conditional reasoning and the dual-source model. *Front. Psychol.* 5:316

Skovgaard-Olsen N, Collins P, Krzyżanowska K, Hahn U, Klauer KC. 2019. Cancellation, negation, and rejection. *Cogn. Psychol.* 108:42–71

Skovgaard-Olsen N, Singmann H, Klauer KC. 2017. Relevance and reason relations. *Cogn. Sci.* 41:1202–15

Sloman SA, Fernbach P. 2017. *The Knowledge Illusion: Why We Never Think Alone*. New York: Riverhead

Sperber D, Wilson D. 1986. *Relevance: Communication and Cognition*. Oxford, UK: Basil Blackwell

Stanovich KE. 2011. *Rationality and the Reflective Mind*. Oxford, UK: Oxford Univ. Press

Stenning K, Cox R. 2006. Reconnecting interpretation to reasoning through individual differences. *Q. J. Exp. Psychol.* 59:1454–83

Stephens RG, Dunn JC, Hayes BK. 2018. Are there two processes in reasoning? The dimensionality of inductive and deductive inferences. *Psychol. Rev.* 125:218–44

Stevenson RJ, Over DE. 1995. Deduction from uncertain premises. *Q. J. Exp. Psychol. Hum. Exp. Psychol.* 48A:613–43

Stevenson RJ, Over DE. 2001. Reasoning from uncertain premises: effects of expertise and conversational context. *Think. Reason.* 7:367–90

Tomasello M. 2014. *A Natural History of Human Thinking*. Cambridge, MA: Harvard Univ. Press

Trouche E, Johansson P, Hall L, Mercier H. 2016. The selective laziness of reasoning. *Cogn. Sci.* 40:2122–36

Trouche E, Sander E, Mercier H. 2014. Arguments, more than confidence, explain the good performance of reasoning groups. *J. Exp. Psychol. Gen.* 143:1958–71

Van Eemeren FH, Grootendorst R. 2004. *A Systematic Theory of Argumentation: The Pragma-Dialectical Approach*. Cambridge, UK: Cambridge Univ. Press

Vidal M, Baratgin J. 2017. A psychological study of unconnected conditionals. *J. Cogn. Psychol.* 29:769–81

Vul E, Goodman N, Griffiths TL, Tenenbaum JB. 2014. One and done? Optimal decisions from very few samples. *Cogn. Sci.* 38:599–637

Walton DN, Reed C, Macagno F. 2008. *Argumentation Schemes*. Cambridge, UK: Cambridge Univ. Press

Wason PC. 1968. Reasoning about a rule. *Q. J. Exp. Psychol.* 20:273–81

Wolf AG, Rieger S, Knauff M. 2012. The effects of source trustworthiness and inference type on human belief revision. *Think. Reason.* 18:417–40

Zhao J, Osherson D. 2010. Updating beliefs in light of uncertain evidence: descriptive assessment of Jeffrey's rule. *Think. Reason.* 16:288–307