
THE REASONER

VOLUME 2, NUMBER 5
MAY 2008

www.thereasoner.org

ISSN 1757-0522

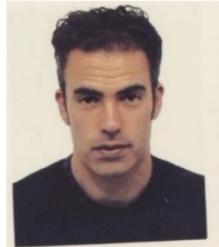
CONTENTS

§1	Editorial	2
§2	Features	3
§3	News	10
§4	Introducing ...	12
§5	Events	16
§6	Jobs	22
§7	Courses and Studentships	23

EDITORIAL

It's a pleasure to be a guest editor for this month. Having spent most of my academic life moving around disciplines, I am a firm believer in interdisciplinary approaches. How else can we study reasoning, without delving into questions of logic, philosophy, AI, statistics, and psychology? Not that we need to address every aspect to make progress, but the bigger picture surely draws on all these areas (and many more). Indeed from the perspective of a cognitive psychologist (what other kinds are there?), it is impossible to explore human reasoning without engaging with most of these disciplines. To understand how people actually reason we need to understand how they *should* reason (even if this is only the first step in the analysis). This throws us into the world of logic, probability theory, Bayesian networks etc. We also want to construct computational models that capture the reasoning process; and understand human reasoning in the broader social and evolutionary context. We want to understand the brain processes that underpin reasoning, the influence of emotions and mood, and so on. The list is endless. It's no wonder that I don't get enough empirical studies done.

The person I've chosen to interview is a perfect exemplar of the interdisciplinary approach. [Nick Chater](#) has pioneered the application of mathematical and computational ideas to psychology, with a healthy respect for attendant philosophical issues (e.g., problem of inductive inference, confirmation, simplicity etc.). Together we run an [MSc in Cognitive & Decision Sciences](#) at UCL. This program integrates a wide range of disciplines and methodologies, with the guiding assumption that human cognition and choice are computational processes, implemented in neural hardware.



[David Lagnado](#)

Psychology, University College London

Interview with Nick Chater

Nick Chater is Professor of Cognitive and Decision Sciences at UCL. He is interested in building and testing formal theories of cognition, often using probabilistic methods and related concepts from information theory. He is co-editor, with Mike Oaksford, of *The Probabilistic Mind* (OUP, 2008), a follow-up to their earlier edited volume *Rational Models of Cognition* (OUP, 1998); and Oaksford and Chater have also written a recent monograph, *Bayesian Rationality* (OUP, 2007), which applies probabilistic methods to resolving a wide range of phenomena in the psychology of human reasoning which appear puzzling from a purely logical point of view.



David Lagnado: You are interested in the ‘rational analysis’ of cognition. What does this approach involve?

Nick Chater: Rational analysis was introduced in psychology by John Anderson in the early 1990s. The idea is to build a formal model of relevant aspects of the environment, the task, and any computational restrictions that must be enforced, and then work out the optimal, or rational, way for the cognitive system to function, under these constraints. This last calculation often involves Bayesian inference. The question is: how far does actual behaviour follow the optimal model? The assumptions of the optimal model can then be iteratively adjusted, to obtain a better fit with the data. This style of explanation is closely related to rational choice explanation (but here computational constraints are typically not considered) and David Marr’s important concept of the computational level of description.

DL: Tell us about the recent probabilistic / Bayesian revolution in approaches to modeling cognition.

NC: Most interesting cognitive problems involve reasoning about uncertainty. For example, perception involves inferring the state of the external world—but there are many ‘worlds’ that yield the same sensory input. So perceptual inference is inherently uncertain. Moreover, each aspect of the perceptual scene is inherently ambiguous (not merely the person, but the face, the eyes, and even the depths, edges or colours in the underlying image). A natural way to model such problems is to measure uncertainty using the probability scale $[0,1]$ —this is to adopt a subjectivist,

or Bayesian, approach to perception; and then a variety of well-known arguments imply that the laws of probability should serve as an appropriate consistency requirement. Many modern computational approaches to vision, language processing, learning, and common-sense reasoning extensively use probabilistic methods. The approach is now also widespread throughout cognitive science and computational neuroscience.

DL: How does this square with the claims by Kahneman and Tversky that people are poor at probabilistic reasoning, and use short-cut heuristics?

NC: Probabilistic models are good at explaining the patterns of human reasoning, perception and so on. But this does not mean that people can explicitly reason about probability. Similarly, the fact that Fourier analysis and wavelet transforms provide a model of early vision is compatible with the fact that most of us have no idea how to explicitly compute the relevant calculations.

DL: What about more recent claims by Gigerenzer and colleagues that people use simple heuristics that are 'ecologically rational'?

NC: One resolution may be that probabilistic analysis is required to understand why certain shortcuts work reliably, given the structure of the environment.

DL: If the mind is probabilistic, what room is there for all-or-none type beliefs and qualitative reasoning?

NC: Indeed, explicit reasoning typically proceeds with all-or-none beliefs; and the main contribution of probability theory may actually be to clarify the appropriate types of qualitative reasoning. Probability theory is really more about the structure of reasoning (roughly, what confirms what) than it is about numbers.

DL: How does the probabilistic mind accommodate the important role of causal representations and causal reasoning? And which is more basic—probabilistic or causal representations? (This is perhaps too simplistic, but Pearl 2000 definitely argues for the primacy of causality).

NC: I think Pearl is absolutely right that causal representations are fundamental. Only with a representation of the causal structure of a situation, can we reason about counterfactuals, or the likely consequences of action. If we understand the causal structure of a machine, for example, we can use it, repair it, break it, or perhaps even convert it to a new use. Similarly, we can reason about what will probably happen in each of these cases. But the notion of causality cannot be reconstructed in purely probabilistic terms.

DL: How important is action to cognition?

NC: Action is probably important in reasoning. If I see two correlated events, A and B, I may not know which caused the other, or whether there was a third

cause C. But if A is the direct result of my own action, then it was not caused by B; and it is natural to suppose that A caused B (it is possible, but very unlikely, that B spontaneously occurs, perhaps caused by C or some other hidden event, just when and only when I bring about A). But the important part of this is really that I know that a particular causal link (between B and A) is not operative. The same information may be available if I see someone else acting to produce A; or if A is caused by what I take to be an unequivocally external event (like a lightning strike); or, of course, if A occurs before B. Action is just one of many clues to causal structure. Action is also important, of course, because it allows us to flexibly collect new information—e.g., I may move around to get a better view; or act to make a particular observation or conduct a crucial experiment.

The Admissibility of Evidence about Previous Convictions in Court I: Setting the Problem

The admissibility of evidence about previous convictions in court has been a contentious issue for a long time. Traditionally, the admissibility of this evidence has been restricted under English law. However, the Criminal Justice Act 2003 (CJA)'s extensive reform includes a more sympathetic approach to its admissibility. This change in the legal position has brought back into focus the question: what, if anything, is wrong with evidence of previous convictions? This paper critically evaluates some of the common responses to this evidence. A sequel, to be published in the next issue of *The Reasoner*, will offer insight as to how this question should be answered.

A common intuitive reaction to previous convictions is that this evidence is irrelevant to the case at hand. However, 'relevance' as used in evidence law is a rather technical term: 'relevant ... evidence is evidence which makes the matter which requires proof more or less probable' (Lord Simon in *DPP v Killbourne* [1973] AC 729, p. 756). To deny that previous convictions are relevant, one has to show that the probability of the accused committing the offence remains completely unaffected after the previous convictions evidence is introduced. Yet considering this evidence to be *irrelevant* is an affront to common sense. You walk in a dark street and suspect that the person who is following you might rob you. Would the probability that he might rob you remain exactly the same if you knew that he had several previous convictions for robbery? (For a comprehensive criticism of the irrelevance argument, see Mike Redmayne 2002: 'The Relevance of Bad Character', Cambridge Law Journal, 684-714).

Of course, it is possible to quibble with the definition of relevance which is used in law or with the interpretation of probability which underlies the claim above. However, in addition to prescribing alternative concepts to those used by the law, one has to be careful not to provide a rationale which would object not only to previous convictions but to *all* evidence. For example, Stein holds that only evidence which is *specific to the particular case* is relevant, and since the applicability of previous convictions to the specific case cannot be tested, they should be excluded (Alex Stein 2005: *Foundations of Evidence Law*, OUP, 183-187). However, this claim denies the legal fact-finder the use of *any* type of evidence. It has long been observed by legal scholars that any inferential step from evidence to the specific case requires a generalisation. To take Stein's own example (2005: 89), if the defence's witness happens to be the accused's brother, then inferring that this fact makes the witness less credible requires the generalisation that interested witnesses tend to be less credible than uninterested witnesses. This generalisation (and any other) is not specific to the case at hand (and testing its applicability to the specific case is not different from previous convictions). Should we exclude the evidence that the witness is the accused's brother just because it requires the fact-finder to use a non-specific generalisation? If so, how can the court do *any* factual finding? If not, what is the difference between this evidence and previous convictions?

Another common objection is that previous convictions are prejudicial because juries *overestimate* their probative value. Yet, it is insufficient to show that juries ascribe this evidence high probative value; one needs to show that the weight given is more than the weight *deserved*. Ample empirical evidence about recidivism shows that having previous convictions is a weighty indicator for future convictions. For example, in the United States, of the 272,111 people released from prisons in 15 states in 1994, approximately 46.9% (!) were rearrested and convicted of a felony (see [US Department of Justice](#)). Not only does it show that previous convictions are relevant and probative, this example also illustrates how hard it is to establish that previous convictions receive *more weight than they deserve* (see also Redmayne).

A more promising type of objection to previous convictions relies on moral considerations. For example, Wasserman argues that using previous convictions evidence against the individual fails to respect his autonomy (David Wasserman 1992: 'The Morality of Statistical Proof and the Risk of Mistaken Liability', 935-976). According to Wasserman, to respect the accused as an autonomous individual, the law has to assume that he is 'free to determine and *alter his conduct at each moment*', Wasserman (1992: 943, emphasis added). Previous convictions evidence is therefore problematic because it makes the accused's liability depend upon 'the

dead hand of his own past', Wasserman (1992: 956-957).

Although this seems to be a move in the right direction, it is still unclear exactly how using previous convictions evidence fails to respect the individual's autonomy. One might take Wasserman to be arguing that the accused's past decisions should not be used as evidence about his present decisions. But this cannot be true. Otherwise, one should object to *any* evidence about the accused's actions before the crime which are by themselves insufficient to constitute an offence (evidence about motive, planning, etc). Furthermore, the past misconduct for which the accused was previously convicted was under the accused's control (putting aside the question of whether this control is based on merely subjective experience or on objective availability of alternate possibilities). If control has anything to do with autonomy, then perhaps using previous convictions evidence actually shows *more* respect to the individual's autonomy. This is because it confronts the individual with the full consequences of his past choices, one of which is that he now belongs to a group which is more likely to re-offend (see the recidivism argument above). Regardless of whether one accepts these objections to Wasserman's account (and I do not), this account should be further developed to explain how admitting previous convictions conflicts with individual autonomy.

The problem of previous convictions is therefore still to be resolved. It is necessary to explain (or explain away) the intuitive objection to this evidence and to evaluate whether or not there is any justification to exclude it from court.

Amit Pundik
Law, Cambridge

A note on probabilistic logics and probabilistic networks

In a classical logic we are typically faced with the following kind of question: do some given premisses $\varphi_1, \dots, \varphi_n$ entail a given conclusion ψ ? This question can be written

$$\varphi_1, \dots, \varphi_n \vDash \psi?$$

Here $\varphi_1, \dots, \varphi_n, \psi$ are premisses of some formal language, such as a propositional language or a predicate language. \vDash is an *entailment* relation: the entailment holds if all models of the premisses also satisfy the conclusion, where the logic provides some suitable notion of 'model' and 'satisfy'. Proof theory is normally invoked to answer a question of this form: one tries to prove the conclusion from the premisses

in a finite sequence of steps, where at each step one invokes an axiom or applies a rule of inference.

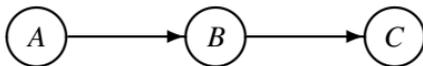
Probabilistic logics come in various guises but we shall look at logics where probabilities attach to sentences. Thus $\forall x(Ux \rightarrow Vx)^{0.8}$ says that the probability that all U s are V s is 0.8. Probabilistic logics have great potential in any application in which logical or structural constraints operate, but where they only operate in a certain proportion of cases or where the constraints are uncertain—e.g., in inferring meaning in natural language, predicting protein folding in biology or modelling scientific theory change.

In a probabilistic logic, the fundamental question takes a different form to that of classical logic. While we might have premisses of the form $\varphi_1^{X_1}, \dots, \varphi_n^{X_n}$ where X_1, \dots, X_n are probabilities or sets of probabilities, it is rare that we are presented with a conclusion of the form ψ^Y and asked whether the conclusion follows from the premisses. More typically, there is a conclusion proposition ψ of interest and we want to know what probability or set of probabilities Y to attach to ψ . Thus the fundamental question of probabilistic logic can be written

$$\varphi_1^{X_1}, \dots, \varphi_n^{X_n} \approx \psi^?$$

Since this question differs from that of classical logic, one might anticipate that the means to solve the question differ too. In fact, determining Y is essentially a question about probability, so methods of probabilistic inference are more appropriate than the standard notion of proof. In (2008: *Probabilistic logics and probabilistic networks*, available [here](#)) [Rolf Haenni](#), [Jan-Willem Romeijn](#), [Gregory Wheeler](#) and I explore the use of *probabilistic networks* to answer this question.

A *Bayesian network*—the simplest kind of probabilistic network—consists of a directed acyclic graph on a finite set of variables, together with the probability distribution of each variable conditional on its parents in the graph. For example, given propositional variables A , B and C , the following constitutes a Bayesian network:



$P(A) = 0.7$	$P(B A) = 0.2$ $P(B \neg A) = 0.1$	$P(C B) = 0.9$ $P(C \neg B) = 0.4$
--------------	---------------------------------------	---------------------------------------

By assuming what is called the *Markov Condition*, which says that each variable is probabilistically independent of its non-descendants in the graph conditional on

its parents, a Bayesian network suffices to determine the joint probability distribution over the set of variables. In our example the Markov condition says that C is independent of A conditional on B . Probabilities over the whole set of variables are multiples of corresponding conditional probabilities: $P(A \wedge \neg B \wedge C) = P(C|\neg B)P(\neg B|A)P(A)$.

A *credal network* is just like a Bayesian network except that the conditional probabilities are only identified to within closed intervals. Thus the above graph together with the following constraints determine a credal network:

$P(A) \in [0.7, 0.8]$	$P(B A) = 0.2$ $P(B \neg A) \in [0.1, 1]$	
<table border="1" style="margin: 0 auto;"> <tr> <td style="padding: 5px;">$P(C B) \in [0.9, 1]$ $P(C \neg B) \in [0.4, 0.45]$</td> </tr> </table>		$P(C B) \in [0.9, 1]$ $P(C \neg B) \in [0.4, 0.45]$
$P(C B) \in [0.9, 1]$ $P(C \neg B) \in [0.4, 0.45]$		

While a Bayesian network represents a single probability function, a credal network represents a convex set of probability functions. A wide variety of algorithms have been developed for constructing probabilistic networks and for calculating probabilities from them. The use of probabilistic networks can greatly reduce the computational burden of probabilistic inference: broadly speaking, the sparser the graph, the quicker it is to draw inferences.

In the context of probabilistic logic it turns out that for a range of natural semantics the X_1, \dots, X_n are normally probabilities or intervals of probabilities, and consequently the premisses determine a convex set of probability functions. A probabilistic network can be used to represent that set of functions and to infer an appropriate set Y of probabilities to attach to the conclusion sentence. While the probabilistic network itself depends on the chosen semantics, the machinery for calculating Y does not—see our (2008: §8.2).

A surprisingly broad range of approaches to probabilistic inference can be invoked to provide semantics for probabilistic logics. Under the *standard probabilistic semantics* a probability function P satisfies φ^X iff $P(\varphi) \in X$; premisses entail a conclusion iff all probability functions that satisfy the premisses also satisfy the conclusion. According to *probabilistic argumentation*, Y is the probability of worlds for which the premisses *force* the conclusion sentence ψ to be true. With *evidential probability*, the φ_i include statistical statements, ψ is inferred by certain rules for manipulating reference classes, and Y quantifies the level of risk associated with this inference. *Classical statistics* can also be used to provide a semantics since

probabilistic argumentation or evidential probability can capture fiducial probability. According to *Bayesian statistical inference* the premisses contain information about prior probabilities and likelihoods, and the entailment holds if the conclusion follows by Bayes' theorem. With *objective Bayesian epistemology*, the entailment holds if any agent with evidence characterised by the premisses should believe ψ to degree within Y .

In sum, probabilistic logics admit a range of natural semantics and probabilistic networks can be used to answer the queries that such a logic faces. The interested reader is urged to consult our (2008) for more details.

Jon Williamson
Philosophy, Kent

§3 NEWS

Interval / Probabilistic Uncertainty and Non-Classical Logic, 25–28 March

Most successful applications of modern science and engineering, from discovering the human genome to predicting weather to controlling space missions, involve processing large amounts of data and large knowledge bases. The ability of computers to perform fast data and knowledge processing is based on the hardware support for super-fast elementary computer operations, such as performing arithmetic operations with (exactly known) numbers and performing logical operations with binary ('true'-'false') logical values.

In practice, experts are usually not 100% certain about the statements included in the knowledge bases, we therefore need to go beyond exact data as well as truth values and such operations from the traditional 2-value logic. A natural way to describe this uncertainty is to use non-classical logics (probabilistic, modal, fuzzy, etc.).

Also, in practical applications, measurements are never 100% accurate. It is therefore necessary to find out how this input inaccuracy (uncertainty) affects the results of data processing. Sometimes, we know the corresponding probability distribution; sometimes, we only know the upper bounds on the measurement error—

which leads to interval bounds on the (unknown) actual value. At present, there are well-developed techniques proposed in the literature for handling situations in which either we know the exact probability distribution for the measurement error or we have no information about the probabilities but only know the upper bound on the measurement error. Particularly, in real life, we also often encounter intermediate situations in which we have partial information about the probabilities frequently described in interval-related terms. To handle such a situation, it is desirable to combine probabilistic and interval approaches to uncertainty. Several formalisms have been developed for such combination as imprecise probabilities, Dempster-Shafer approach, rough-set related approaches, and many others.

Furthermore, in solving many practical problems, we sometime need to simultaneously process both measurement data and expert knowledge which are inaccurate and uncertain in nature as mentioned above. Then it is desirable to combine uncertainty analysis with non-classical logic. Such a combination would also play an important role in developing intelligent systems, where the former would serve for modelling and representing knowledge, and the latter would serve as the foundation for reasoning and supporting decision-making.

The above-mentioned combinations were the main objectives of the [International Workshop on Interval/Probabilistic Uncertainty and Non-classical Logics \(UncLog08\)](#), which was successfully held at Japan Advanced Institute of Science and Technology during March 25-28, 2008, and of its proceedings. The workshop brought together over thirty participants, including world renowned researchers and young active scientists working on the related fields, from countries as diverse as China, Czech Republic, France, Germany, Japan, Malaysia, Portugal, Sweden, Switzerland, Thailand, United States, United Kingdom, and Vietnam. It provided an exciting forum for the discussion and exchange of research results and ideas on the specified topics among participants.

We hope that this workshop will lead to a boost in the much-needed collaboration between the uncertainty analysis and non-classical logic communities, and thus, to better processing of uncertainty.

[Van Nam Huynh](#)

Japan Advanced Institute of Science and Technology

[Vladik Kreinovich](#)

Computer Science, University of Texas at El Paso

Calls for Papers

INFORMATION FUSION: Information Fusion in Public Health Informatics and Surveillance, special issue of Information Fusion, deadline 30 May.

APPLICATIONS AND METHODOLOGIES FOR PLANNING AND SCHEDULING: Special issue of Journal of Scheduling, deadline 15 June.

CAUSALITY AND PROBABILITY IN THE SCIENCES

Deadline 1 July

PROBABILISTIC MODELS FOR IMAGE UNDERSTANDING: Special Issue of the International Journal of Computer Vision, deadline 21 July.

KYBURG: Special issue of Synthese commemorating Henry E. Kyburg, Jr, deadline 30 July.

PROBABILISTIC GRAPHICAL MODELS IN COMPUTER VISION: Special issue of IEEE Transactions on Pattern Analysis and Machine Intelligence, deadline 16 August.

CONDITIONALS AND RANKING FUNCTIONS: Special issue of Erkenntnis, franz.huber@uni-konstanz.de, deadline 31 August.

DEPENDENCE ISSUES IN KNOWLEDGE-BASED SYSTEMS: Special Issue of International Journal of Approximate Reasoning, deadline 15 September.

§4

INTRODUCING ...

In this section we introduce a selection of key terms, texts and authors connected with reasoning. Entries will be collected in a volume *Key Terms in Logic*, to be published by Continuum. If you would like to contribute, please [click here](#) for more information. If you have feedback concerning any of the items printed here, please email thereasoner@kent.ac.uk with your comments.

Analogy

An analogy is a logical function and a reasoning process. As a logical function, analogy is a relation between two ordered pairs: in ‘A is to B as C is to D’, the particular elements are different but the functional relation is identical. As a reasoning process, an analogy is a structure-preserving mapping between two conceptual domains: a ‘source’ and a ‘target’, the first one being familiar, the second unfamiliar.

Properties of items in the source domain are applied to items in the target domain. Analogy is thought to be important in abduction.

Benoit Hardy-Vallée
Philosophy, Toronto

Computability

A function is computable when it is calculable in a definite (i.e., algorithmic) way in a finite number of steps. A set is computable when its characteristic function (the function f such that $f(x) = 1$ when x belongs to the set) is computable. Several precise definitions have been proposed in order to explicate this somewhat vague notion, as a result of the works of Church, Gödel, Kleene, and Turing. Computable functions have been defined in terms of recursive functions, Turing machines, and Lambda calculus. All these definitions turned out to be exactly coextensive. Recognizing these equivalences, the Church-Turing thesis states that a function is computable when it is calculable by a Turing machine (or, equivalently, when it is a recursive function). The Church-Turing thesis offers a rigorous explanation of computable functions; the thesis, which is generally accepted, can be neither proved nor refuted because it is an explication of an informal notion. A consequence of the Church-Turing thesis is that computable functions are those functions calculable by computers, which are physical realizations of universal Turing machine.

There is a strong connection between computability and decidability: a theory is decidable when the set of its valid formulae is computable. In 1936 Turing, answering in the negative the decision problem for predicate logic posed by Hilbert, constructed a non computable set of natural numbers by means of a diagonal argument similar to those previously employed by Cantor and Gödel. In this way Turing proved the existence of clearly defined non computable functions.

Mauro Murzi

Fallacy

A fallacy is an argument which, though plausible, is defective. 'Argument' is here used in the broad sense of a contribution to a dialogue, in which reasons are advanced. There are rules for the proper conduct of an argument. To commit a fallacy

is to perpetrate (accidentally or deliberately) a subtle violation of one or other of those rules.

There are many types of argumentative misconduct and, correspondingly, many types of fallacy. For example, an opponent should not be bullied into accepting anything. To break this rule by insinuating some kind of a threat into one's argument is to commit the fallacy *Argumentum ad Baculum*—an appeal to force. Appeal to an authority (especially when that authority is not an authority on the matter under discussion) is the fallacy *Argumentum ad Verecundiam*. Appealing to popular opinion (*Argumentum ad Populum*) is fallacious, not so much because public opinion is frequently mistaken, as because citing the widely held opinion in favour of a view does not amount to supplying a solid *argument* showing that view to be true.

An obvious rule of argumentative conduct is that we should express ourselves in such a way that we are not likely to be misunderstood. We commit the fallacy of *Equivocation* when one or more of the words used in the argument is ambiguous, leaving ourselves open to being interpreted in more than one way. Sometimes a sentence can be susceptible of two or more interpretations not because any of the words in it are ambiguous but because it may be parsed in two or more ways. Example: 'Fighting kangaroos can be fun'. Argumentative error that arises from a confusing of such interpretations is known as the fallacy of *Amphiboly*.

Another rule for good argument is that the premises should be relevant to establishing the conclusion. Some fallacies arise when premises that might seem to be relevant or adequate to establishing the conclusion are in fact irrelevant or inadequate:

Post hoc ergo propter hoc (After this, therefore on account of this). It is a fallacy to think that if one event occurs after another, the earlier accounts for, or causes, the later. It may be true that, after schools started employing educational psychologists, the rate of teenage drug abuse and crime rose. But it does not follow that employing educational psychologists was responsible for this rise.

Petitio Principii (Also known as Begging the Question or Circular Argumentation). This fallacy occurs when one of the reasons given in the course of an argument is identical to, or just a linguistic variant on, the very conclusion one is trying to draw. The simplest example of this is when premise and conclusion are equivalent, e.g., 'It is best to have government by the people because democracy is the best form of government'. But much more complicated cases are usual. (Note that there is a widespread but mistaken tendency to think that 'to beg the question' means 'to raise the question'. To beg the question is to argue for a particular answer in a way that presupposes the truth of that answer.)

Hasty generalization. The mistake of drawing a general conclusion from a small number of possibly atypical instances. That John Stuart Mill and Bertrand Russell profited greatly from a private education does not imply that all children would benefit from being educated at home.

Straw Man. In this fallacy your argument is relevant to establishing a conclusion subtly different from the conclusion you are purporting to establish. Suppose the bus company wants to raise its fares by 20%. That will hit you badly in the pocket, and you write a long letter to the company arguing that a 20% increase in the cost of living would be economically ruinous. But neither the bus company nor anyone else is proposing a 20% increase in the general cost of living. So your letter is aiming at the wrong target, one of your own contrivance, a straw man.

Argumentum ad Ignorantiam. The fallacy of arguing from the fact that we don't know whether something is true to the conclusion that it is false. It may be true, for example, that none of the research that has been done on the subject of pornography has established that it does harm to women. So we don't know whether pornography does harm to women. But it would be fallacious to conclude that pornography does not harm women.

Denying the Antecedent. Consider 'If Beavis continues to smoke, he will die within three years. Beavis does not continue to smoke. Therefore he will not die within three years.' This might seem to be a valid argument, but it is not—the premises could be true while the conclusion is false. This is a formal fallacy—any argument of the same form (If A then B. Not-A. Therefore Not-B) is fallacious. Another formal fallacy is *Affirming the Consequent* (If A then B. B. Therefore A.)

Logicians study fallacies (of which the above is but a small sample) because it is useful to identify and categorize the ways in which we ought *not* to reason, and the classification of fallacies is usually in terms of the type of logical error they exhibit. However, psychological experiments have shown that, when reasoning about factual matters, humans are subject to diverse biases and misconceptions. We could, perhaps more fruitfully, classify fallacies according to their psychological roots. It should be possible to find deep evolutionary explanations for the various sorts of argumentative errors we are prone to commit. Revealing the ways we go wrong is very revealing of the kind of animal we are.

In his excellent compendium of fallacies *How to Win Every Argument* (Continuum, 2006), Madsen Pirie recommends arming oneself with the impressive Latin names of fallacies before doing argumentative battle: 'When an opponent is accused of perpetrating something with a Latin name it sounds as if he is suffering from a rare tropical disease. It has the added effect of making the accuser seem both erudite

§5 EVENTS

MAY

SBIES: Seminar on Bayesian Inference in Econometrics and Statistics, University of Chicago Graduate School of Business Gleacher Center, 2–3 May.

TML: Workshop on Teaching Machine Learning, Saint-Etienne, 5–7 May.

PRAGMATISM AND NATURALISM: Workshop, Tilburg Center for Logic and Philosophy of Science, 7–9 May.

SIG16: 3rd Biennial Meeting of the EARLI-Special Interest Group 16—Metacognition, Ioannina, Greece, 8–10 May.

CLE, EBL & SLALM: 30th Anniversary of the Centre for Logic, Epistemology and the History of Science (CLE), UNICAMP, 15th Brazilian Logic Conference, and 14th Latin-American Symposium on Mathematical Logic, Paraty, Brazil, 11–17 May.

ARGMAS: Fifth International Workshop on Argumentation in Multi-Agent Systems, Estoril, Portugal, 12–13 May.

INTERVAL PROBABILITY: Workshop on Principles and Methods of Statistical Inference with Interval Probability, Durham, 12–16 May.

DL: 21st International Workshop on Description Logics, Dresden, 13–16 May.

FEW: Fifth Annual Formal Epistemology Workshop, Madison, Wisconsin, 14–18 May.

UR: Special Track on Uncertain Reasoning, 21st International Florida Artificial Intelligence Research Society Conference, Coconut Grove, Florida, 15–17 May.

AI PLANNING AND SCHEDULING: A Special Track at the 21st International FLAIRS Conference, Coconut Grove, Florida, 15–17 May.

RSKT: Rough Sets and Knowledge Technology, Chengdu, 17–19 May.

MANYVAL: Applications of Topological Dualities to Measure Theory in Algebraic Many-Valued Logic, Milan, 19–21 May.

NAFIPS: North American Fuzzy Information Processing Society Annual Conference, Rockefeller University, New York, 19–22 May.

ISMIS: The Seventeenth International Symposium on Methodologies for Intelligent Systems, York University, Toronto, Canada, 20–23 May.

WCB: Workshop on Constraint Based Methods for Bioinformatics, Paris, 22 May.

APPROXIMATE INFERENCE: PASCAL 2008 Workshop on Approximate Inference in Stochastic Processes and Dynamical Systems, Cumberland Lodge, 27–29 May.

COMMA: Second International Conference on Computational Models of Argument, Toulouse, 28–30 May.

AI: 21st Canadian Conference on Artificial Intelligence, Windsor, Ontario, 28–30 May.

EXPRESSIONS OF ANALOGY: Faculty of Social and Human Sciences, New University of Lisbon, 29–31 May.

JUNE

AREA: International Workshop on Advancing Reasoning on the Web: Scalability and Commonsense, Tenerife, 1 June.

WCCI: IEEE World Congress on Computational Intelligence, Hong Kong, 1–6 June.

ULTRAMATH: Applications of Ultrafilters and Ultraproducts in Mathematics, Pisa, 1–7 June.

META-ANALYSIS: Synthesis and Appraisal of Multiple Sources of Empirical Evidence, Statistical and Applied Mathematical Sciences Institute, North Carolina, 2–13 June.

CSHPS: Canadian Society for History and Philosophy of Science, University of British Columbia, Vancouver, 3–5 June.

CIE: Computability in Europe 2008: Logic and Theory of Algorithms, University of Athens, Athens, 15–20 June.

MATHEMATICAL PRACTICES: Seville, 16–17 June.

IIS: Intelligent Information Systems, Zakopane, Poland, 16–18 June.

DM: SIAM Conference on Discrete Mathematics, University of Vermont, Burlington, Vermont, 16–19 June.

LOGICA: Hejnice, Czech Republic, 16–20 June.

IEA-AIE: 21st International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, Wroclaw, Poland, 18–20 June.

HOPOS: Seventh Congress of the International Society for the History of Philosophy of Science, Vancouver, Canada, 18–21 June.

HDM: Multivariate statistical modelling and high dimensional data mining, Kayseri, Turkey, 19–23 June.

EPISTEME: Law and Evidence, Dartmouth College, 20–21 June.

IS SCIENCE INCONSISTENT?: History and Philosophy of Science, University of Leeds, 21 June.

IPMU: Information Processing and Management of Uncertainty in Knowledge-Based Systems, Malaga, Spain, 22–27 June.

MED: 16th Mediterranean Conference on Control and Automation, Ajaccio, Corsica, 25–27 June.

ESPP: European Society for Philosophy and Psychology, Utrecht, 26–28 June.

PHILOSOPHY OF PROBABILITY: Graduate Conference, London School of Economics, 27–28 June.

DGL: Second Workshop in Decisions, Games and Logic, Institute for Logic, Language and Computation, Amsterdam, 30 June – 2 July.

EWRL: European Workshop on Reinforcement Learning, INRIA, Lille, 30 June – 3 July.

JULY

WoLLIC: 15th Workshop on Logic, Language, Information and Computation, Edinburgh, 1–4 July.

LOFT: 8th Conference on Logic and the Foundations of Game and Decision Theory, 3–5 July.

LOGIC COLLOQUIUM: Bern, Switzerland, 3–8 July.

ICML: International Conference on Machine Learning, Helsinki, 5–9 July.

SMT: 6th International Workshop on Satisfiability Modulo Theories, Princeton, 7–8 July.

COMPUTATION AND COGNITIVE SCIENCE: King's College, Cambridge, 7–8 July.

NEGATION AND DENIAL: Philosophy Centre, University of Lisbon, 7–8 July.

CAV: 20th International Conference on Computer Aided Verification, Princeton, 7–14 July.

INDUCTION: Historical and Contemporary Approaches, 5th Ghentian Conference in the Philosophy of Science, Centre for Logic and Philosophy of Science, Ghent, 8–10 July.

BAYESIAN MODELLING: 6th Bayesian Modelling Applications Workshop, Helsinki, 9 July.

EVALUATING AND DISSEMINATING PROBABILISTIC REASONING SYSTEMS: Helsinki, 9 July.

UAI: Uncertainty in Artificial Intelligence, Helsinki, 9–12 July.

COLT: Conference on Learning Theory, Helsinki, 9–12 July.

CLASSICAL LOGIC AND COMPUTATION: Reykjavik, 13 July.

WCP4: Fourth World Congress of Paraconsistency, Melbourne, 13–18 July.

BPR: The 1st International Workshop on Bit-Precise Reasoning, Princeton, 14 July.

ITSL: Information Theory and Statistical Learning, Las Vegas, 14–15 July.

IKE: International Conference on Information and Knowledge Engineering, Las Vegas, 14–17 July.

DMIN: International Conference on Data Mining, Las Vegas, 14–17 July.

NORMAS: 3rd International Workshop on Normative Multiagent Systems, Luxembourg, 15–16 July.

DEON: 9th International Conference on Deontic Logic in Computer Science, Luxembourg, 15–18 July.

NCPW: 11th Neural Computation and Psychology Workshop, Oxford, 16–18 July.

PROOF THEORY: Workshop on Logic, Foundational Research, and Metamathematics II, WWU Institute for Mathematical Logic, Münster, 18–19 July.

MoChART: Fifth Workshop on Model Checking and Artificial Intelligence, Patras, Greece, 21–22 July.

WIGSK: Inference methods based on graphical structures of knowledge, Patras, Greece, 21–22 July.

ISBA: 9th World Meeting, International Society for Bayesian Analysis, Hamilton Island, Australia, 21–25 July.

INTERDISCIPLINARY SOCIAL SCIENCES: Monash University Centre, Prato, Tuscany, Italy, 22–25 July.

MODEL SELECTION: Current Trends and Challenges in Model Selection and Related Areas, University of Vienna, 24–26 July.

WHAT (GOOD) IS HISTORICAL EPISTEMOLOGY?: Max Planck Institute for the History of Science, Berlin, 24–26 July.

ICHST: XXIIIrd Congress of History of Science and Technology, Budapest, 26–31 July.

ESARM: Workshop on Empirically Successful Automated Reasoning for Mathematics, Birmingham, UK, 26 July – 2 August.

FIRST FORMAL EPISTEMOLOGY FESTIVAL: Conditionals and Ranking Functions, Konstanz, 28–30 July.

AUGUST

LANGUAGE, COMMUNICATION AND COGNITION: University of Brighton, 4–7 August.

ESSLLI: European Summer School in Logic, Language and Information, Freie und Hansestadt Hamburg, Germany, 5–15 August.

BLAST: Boolean Algebra, Lattice Theory, Algebra, Set Theory and Topology, Denver, 6–10 August.

IJCAR: The 4th International Joint Conference on Automated Reasoning, Sydney, 10–15 August.

DEMA: Designed Experiments: Recent Advances in Methods and Applications, Isaac Newton Institute, Cambridge, 11–15 August.

ICT: The Sixth International Conference on Thinking, San Servolo, Venice, 21–23 August.

COMPSTAT: International Conference on Computational Statistics, Porto, Portugal, 24–29 August.

FSKD: The 5th International Conference on Fuzzy Systems and Knowledge Discovery, Jinan, China, 25–27 August.

LSFA: Third Workshop on Logical and Semantic Frameworks, with Applications, Salvador, Bahia, Brazil, 26 August.

LOGICAL PLURALISM: University of Tartu, Estonia, 27–31 August.

NORMATIVITY: Graduate Philosophy Conference on Normativity, Amsterdam, 29–30 August.

SEPTEMBER

IVA: The Eighth International Conference on Intelligent Virtual Agents, Tokyo, 1–3 September.

GRANDEUR OF REASON: Rome, 1–4 September.

10TH ASIAN LOGIC CONFERENCE: Kobe University, Japan, 1–6 September.

COMSOC: 2nd International Workshop on Computational Social Choice, Liverpool, 3–5 September.

KES: 12th International Conference on Knowledge-Based and Intelligent Information & Engineering Systems, Zagreb, 3–5 September.

ICANN: 18th International Conference on Artificial Neural Networks, Prague, 3–6 September.

BLC: British Logic Colloquium, Nottingham, 4–6 September.

NATURALISM: Kazimierz Naturalism Workshop, Kazimierz Dolny, Poland, 6–10 September.

SMPS: Soft Methods for Probability and Statistics, 4th International Conference, Toulouse, 8–10 September.

AIML: Advances in Modal Logic, LORIA, Nancy, France, 9–12 September.

CAUSALITY AND PROBABILITY IN THE SCIENCES

University of Kent, Canterbury UK, 10–12 September

COLLOQUIUM LOGICUM: The biennial meeting of the German Society for Mathematical Logic, Technische Universitaet Darmstadt, 10–12 September.

LOGIC OF CHANGE, CHANGE OF LOGIC: Prague, 10–14 September.

NMR: Twelfth International Workshop on Non-Monotonic Reasoning, Special Session on Foundations of NMR and Uncertainty, Sydney, 13–15 September.

ICAPS: International Conference on Automated Planning and Scheduling, Sydney, 14–18 September.

ECML PKDD: The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Antwerp, Belgium, 15–19 September.

SPATIAL COGNITION: Schloss Reinach, Freiburg, 15–19 September.

CSL: 17th Annual Conference of the European Association for Computer Science Logic, Bertinoro, Italy, 15–20 September.

PGM: The fourth European Workshop on Probabilistic Graphical Models, Aalborg, Denmark, 16–19 September.

KRAMAS: Workshop on Knowledge Representation for Agents and Multi-Agent Systems, Sydney, 16–19 September.

HAIS: 3rd International Workshop on Hybrid Artificial Intelligence Systems, Burgos, Spain, 24–26 September.

OCTOBER

SUM: Second International Conference on Scalable Uncertainty Management, Naples, 1–3 October.

SETN: 5th Hellenic Conference on Artificial Intelligence, Syros, Greece, 2–4 October.

REASON, ACTIVISM, AND CHANGE: University of Windsor, 3–5 October.

ICAI: The 1st International Conference on Advanced Intelligence, Beijing, 19–22 October.

FotFS VII: Bringing together Philosophy and Sociology of Science, Foundations of the Formal Sciences VII, Vrije Universiteit Brussel, 21–24 October.

MICAI: 7th Mexican International Conference on Artificial Intelligence, Mexico City, 27–31 October.

MDAI: Modeling Decisions for Artificial Intelligence, Barcelona, 30–31 October.

NOVEMBER

AUTOMATED SCIENTIFIC DISCOVERY: AAAI Fall Symposium, Arlington, Virginia, 7–9 November.

GAME THEORY: 5th Pan-Pacific Conference in Game Theory, Auckland, 19–21 November.

DECEMBER

ICLP: 24th International Conference on Logic Programming, Udine, Italy, 9–13 December.

ICDM: 8th IEEE International Conference on Data Mining, Pisa, 15, 19 December.

PRICAI: Tenth Pacific Rim International Conference on Artificial Intelligence, Hanoi, Vietnam, 15–19 December.

§6

JOBS

ENGINEERING MATHEMATICS, BRISTOL: Research Assistant, machine learning, statistics or bioinformatics, deadline 12 May.

PHILOSOPHY, ABERDEEN: Lecturer / Senior Lecturer / Reader, deadline 16 May.

THEORETICAL PHILOSOPHY, GHENT: Lecturer or senior lecturer, deadline 23 May.

STATISTICS, QUEENSLAND: Modelling risk in stochastic population networks, Research Fellowship, University of Queensland, deadline 30 May.

STATISTICS, LOUVAIN: 3 Postdoc positions, deadline 1 June.

COURSES AND STUDENTSHIPS

Courses

MA IN REASONING

An interdisciplinary programme at the University of Kent, Canterbury, UK. Core modules on logical, causal, probabilistic, scientific and mathematical reasoning and further modules from Philosophy, Psychology, Computing, Statistics and Law.

EASSS: 10th European Agent Systems Summer School, New University of Lisbon, 5–9 May.

REDUCTIONISM AND EMERGENCE: Tuebingen, 5–9 May.

LOGIC SCHOOL: State University of Campinas, Brazil, 7–9 May.

LOGIC AND FORMAL EPISTEMOLOGY: Summer school for undergraduates, Department of Philosophy, Carnegie Mellon University, Pittsburgh, 9–27 June.

SIPTA: 3rd SIPTA School on Imprecise Probabilities, Montpellier, 2–8 July.

MSc IN COGNITIVE & DECISION SCIENCES

University College London.

PROBABILISTIC CAUSALITY: Central European University, Budapest, 21 July–1 August.

GSSPP: Geneva Summer School in the Philosophy of Physics, 22 July–8 August.

LOGIC PROGRAMMING AND COMPUTATIONAL LOGIC: 3rd International Complog/ALP Summer School, New Mexico State University, 24–27 July.

ESSLLI: European Summer School in Logic, Language and Information, Hamburg, 4–15 August.

MATHEMATICS, ALGORITHMS, AND PROOFS: Summer School, Abdus Salam International Centre for Theoretical Physics, Trieste, 11–29 August.

CAUSALITY STUDY FORTNIGHT

University of Kent, Canterbury UK, 8–19 September

Studentships

STATISTICS, LOUVAIN: 8 PhD positions, deadline 1 June.

[BSPS DOCTORAL SCHOLARSHIP](#): Philosophy of Science, UK, deadline 1 August.

Acknowledgements

The Reasoner is a development of the [proginet](#) academic network supported by the [Leverhulme Trust](#).