# The Reasoner

Volume 8, Number 2
February 2014

www.thereasoner.org
ISSN 1757-0522

## Contents

Editorial 11
Features 11
News 15
What's Hot in ... 17
Events 19
Courses and Programmes 20
Jobs and Studentships 21

## Editorial

It is a great pleasure for me to introduce my second interviewee, John F. Horty, to our readers. Jeff is Professor at the Philosophy Department and Institute for Advanced Computer Studies at the University of Maryland (PhD in Philosophy in 1986, University of Pittsburgh), and Affiliate Professor in the Computer Science Department (also at the University of Maryland). He mainly works in Philosophical Logic, AI, and reasoning broadly construed. It is safe to say that Jeff is one of the most prominent philosophers associated with these areas. Having a background in computer science too, he was amongst the first philosophers to apply methods from computer science to philosophical questions. A lot of work in philosophical logic is done in computer science these days, which makes Jeff also a logical choice for one of the four editors of the *Journal of Philosophical Logic* (2008–present).

I first came across Jeff's work when I was a graduate student at the University of Salzburg. Deontic logic was popular in Salzburg, so I started reading his book *Agency and Deontic Logic* (Oxford University Press, 2001). At the time, I was already familiar with some deontic logic, but compared to what I have read in deontic logic until then, Jeff's book was refreshingly different. Agency and Deontic Logic contains a philosophically and conceptually solid, state-of-the-art deontic logic, and even some applications of it. Rather than dealing primarily with technical issues of his theory, Jeff focuses on conceptual and philosophical questions, with his formal theory being equipped to answer them. Agency and Deontic Logic is not just a book in deontic logic, but can also be considered to be a prime example of formal philosophy in general. Of course, this last property also applies to his more recent books *Frege on Definitions: A Case Study of Semantic Content* (Oxford University Press, 2007) and *Reasons as Defaults* (Oxford University Press, 2012). Two of Jeff's papers have been selected for the *Philosopher's Annual* (1995, 2012), which shows that his work is not only of interest to logicians, but of very general philosophical importance. Obviously, all of this makes Jeff a perfect interviewee for *The Reasoner*, and I am very grateful that he took the time to tell us more about his work.

Albert J.J. Anglberger
MCMP, LMU Munich

## Features

### Interview with John F. Horty

Albert Anglberger: What is your academic background and how did you get drawn to the questions you are dealing with

now?

John F. Horty: I studied classics as an undergraduate, at Oberlin College, and, after a couple of years working in a real job, went to graduate school in philosophy at the University of Pittsburgh. My original intention was to specialize in ancient philosophy, but I was instantly diverted to logic. In fact, I had already been interested in logic in college, and had taken reading courses using some classic texts—Mendelson, even Church. But the impression left by those books was that logic existed as a body of established results, discovered by legendary geniuses living in the past, with only a few open problems, of formidable mathematical difficulty, remaining for future geniuses to solve. What I learned in graduate school was that logic was not just a collection of well-defined problems waiting to be solved, but a living field in which ideas on a range of topics could be explored in a precise way, and to which even a non-genius could usefully contribute. Every logician has a "moment" and mine was proving a cut elimination theorem for a simple relevance logic that I'd made up on my own, in a class with Nuel Belnap.

I should also say that I belonged to, I believe, the last generation of Pittsburgh students to be able to study with both Nuel and Rich Thomason, a remarkable experience, as anyone familiar with both of these individuals will know.

The Pittsburgh philosophy department at that time contained several other extraordinary teachers, but what was really important for me was beginning to work—originally as a research assistant, while still a graduate student, then as a research scientist—in the computer science department at Carnegie Mellon University. This gave me the chance to take some graduate AI courses, to hack LISP code for a living, and to hang out with a group of cheerful computer science students—very different from my brooding philosophy friends—who seemed to frame any issue in terms of algorithms, data structures, search, and complexity. I learned about a whole new set of problems, but also, about a new way to think about the problems I'd been familiar with. Eventually I arrived at the view best expressed, some years later, in Johan van Benthem's wonderful observation, echoing Clausewitz, that artificial intelligence is merely a continuation of philosophy by other means.

While at Carnegie Mellon, in the mid 1980s, I had the opportunity to join Rich Thomason, Dave Touretzky, and Chuck Cross in an NSF project on nonmonotonic inheritance reasoning. This provided my introduction to nonmonotonic logic, during the first or second wave of research in that field. It is a wonderful thing to work in a field when it is young, when ideas are fresh, and when it is still driven by problems rather than theories.

AA: You have been one of the first people to apply ideas from computer science to philosophical questions at great length. What have been the merits of this approach so far and where have fruitful but obvious applications been neglected until now?

JH: Well, there are several areas in which there is no longer any noticeable boundary between work by philosophers and work by computer scientists. One is the general area of causality, causal reasoning, and counterfactual reasoning. Also, learning, belief revision, both qualitative and quantitative, decision theory, and probabilistic reasoning more generally—these are all fruitful areas of interaction. In many traditional areas of philosophical logic, I would say that the center of gravity now lies in computer science. This is especially true of temporal and epistemic logics, areas first explored by philosophers, largely stagnant for many years, revived by important work in computer science, and now attracting attention by philosophers again as well. And of course, philosophers at least since Aristotle have talked about or gestured at defeasibility, but it was not until the work on nonmonotonic reasoning in computer science that we began to get any logical understanding of the concept.

As far as missed opportunities—or opportunities that have so far been missed—I will mention only one. There is a central field of philosophy, sometimes called the theory of practical reasoning, focusing on interactions among the concepts of intentions, goals, beliefs, reasons, and aimed, very broadly, at articulating an account of rational action. There is also, in computer science, a body of research on the overall design, or architecture, of an intelligent agent, again involving the concepts of intentions, goals, beliefs, reasoning, rationality, and action. Although the computer science research has, in fact, been impacted by the philosopher Michael Bratman's work on the ways in which intentions, or plans, can constrain reasoning, I would argue that they have drawn the wrong lessons from this work; and there has been no flow at all in the other direction. This is a shame. There should be joint workshops, volumes. While the problems faced by these two communities are not exactly identical, there is certainly enough overlap that more interaction would benefit everyone.

AA: It's commonly known that nonmonotonic logics are closer to commonsense reasoning than monotonic logics. When nonmonotonic logics came up in computer science in the 1980s, some philosophers immediately realized that these logics could very well serve as a basis for various systems of deontic logic. What are the benefits from going nonmonotonic when modeling deontic notions?

JH: The benefit of going nonmonotonic is that deontic reasoning, like just about all of our reasoning in every area except mathematics, is defeasible, and nonmonotonic logics allow us to capture that. I should note that there are two different ways, not always distinguished, in which nonmonotonic logic can serve as a basis for deontic logic. The first is to preserve the ordinary possible worlds semantics for deontic logic, but to graft this standard deontic logic onto an underlying logic containing defeasible conditional, of some kind. This is the approach found in work by Michael Morreau, and by Nicholas Asher and Dan Bonevac, all using model preference defeasible conditionals, by Don Nute, using his own defeasible logic, and by Henry Prakken, using an argumentation-based conditional. The second is to abandon possible worlds semantics for deontic logic, and to appeal, instead, to techniques of nonmonotonic logic in order to explain how the truth or falsity of ought statements is dependent on other sorts of things—reasons, or imperatives. This is the approach taken in my own work, based on default logic, which generalizes Bas van Fraassen's original proposal for basing oughts on imperatives. This approach is also found in the work on imperatives by Jörg Hansen, and in the input/output logics of David Makinson and Leon van der Torre.

AA: Almost any new system of deontic logic is evaluated

nowadays with respect to how it copes with so-called paradoxes of classic deontic logic. Do you think the importance of these paradoxes is exaggerated, underestimated or just about right?

JH: Exaggerated—if what you mean are things like Ross's paradox, or the Good Samaritan paradox. I view these as entirely pragmatic, like the paradoxes of material implication, and never worry about them at all, really. I do realize, though, that I am in a tiny minority here, and that most people do think that the classic deontic paradoxes are important.

AA: Deontic logic was not just meant to be of theoretical interest alone. Since its beginning, logicians emphasized its applicability to moral and legal reasoning. Considering the vast number of deontic logics developed, their applications are still very rare. What do you think have been the most fruitful applications of deontic logic so far and why are there so few of them?

JH: I disagree with the premise of your question. I actually think that there are surprisingly many applications of, or at least appeals to, deontic logic in ethics. Of course, no one feels that you can just read off an ethical theory from a logic, so there are no philosophical applications in that sense. But just look at the recent work of John Broome, Stephen Finlay, or Mark Schroeder, for instance, and you'll see ethicists working out a sort of semantics for ought statements, subject to various constraints—constraints deriving from ethical theory, linguistic constraints—and then testing plausibility in part by exploring the logical features that their semantic accounts commit them to. This is exactly what deontic logic is about.

You might also wonder about applications of deontic logic in AI, and here the question is more complicated, since it gets tangled up with questions concerning the state of logical AI in general. However, setting this larger question aside, and in an attempt to crowdsource an answer, I've put your query to the computer scientists Leon van der Torre and Jan Broersen. Leon notes the area of regulatory compliance, where compliance needs to be verified, as a "killer app" for deontic logic in computer science, and also mentions the large community working in normative multi-agent systems, where deontic logic is used in reasoning about the relevant norms. Jan points out that, although there are few fielded applications of deontic logics at this moment, things may be different in ten year's time, "when the roads are full of driverless cars, hospitals are full of robots, the sky is full of drones, and multiple accidents and violations have occurred"—then maybe we will devote more effort to "applying logics concerning norms, responsibility, knowledge, and intention."

AA: When we take a look at your more recent work, we see a deviation from deontic logics. It seems that you got more and more interested in developing and applying formal methods to questions of moral philosophy as well as the philosophy of law. In 2011, for example, you published a paper entitled "Rules and Reasons in the Theory of Precedent". Can you briefly tell us what this paper is about?

JH: OK, but one point first. I've never really thought of myself as specializing in deontic logic. I'm interested in reasoning, especially human commonsense reasoning, very broadly. One of the things we humans tend to reason about, and a very interesting thing, is what we ought to do. As a result, I've developed a side interest in what we ought to do, and in deontic logic.

But to the paper: The common law, or "case law", as it's also known, is the fundamental legal system in about a third of the world, and even where it is not officially recognized as the fundamental system, as in civil law countries, you still find common law reasoning, which is a very natural form of reasoning that we all engage in. I started thinking about this kind of reasoning both because I thought it would be a fruitful application domain for nonmonotonic logics, which it is, but also because I've always been interested, ever since reading Aristotle's Ethics, in understanding how we arrive at judgments in the particular, concrete situations before us, especially when all the rules or guidelines in our possession seem to support different judgments—this topic is now discussed in ethics under the heading of moral particularism. So the common law was a natural place to look.

The goal of the paper you mention, which is part of a larger project, is to understand the common law doctrine of precedent, according to which decisions by earlier courts somehow constrain the decisions of later courts, while still allowing those later courts a degree of freedom in responding to fresh circumstances. How does that work—how is this balance between constraint and freedom achieved? Some writers argue that earlier courts lay down rules, but that later courts are allowed to modify those rules in limited ways, registering exceptions as appropriate. What I suggest, instead, is that decisions by earlier courts establish a priority ordering among reasons, and that later courts are constrained only to respect the priorities already established, keeping the overall ordering relation consistent. This view has, I argue, many advantages, and provides a very different picture of what is going on in a common law system. It casts the law as a mechanism though which our individual priority orderings among reasons are combined into a social ordering in a way that is incremental, distributed, and responsive to particular circumstances.

AA: In 2012 you published a new book which might be of special interest—of course, not just because of its title!—to our readers. It's called *Reasons as Defaults* and staying true to your nonmonotonic roots you take a default logic as your starting point in order to model what it means to have a reason to do something. Would you please tell us a bit more about this latest book of yours and also why and how your theory is different from the more traditional ones (e.g., game theoretic and probabilistic ones)?

JH: Yes. If you think about it, much of our ordinary reasoning involves an appeal to, or at least a reflection on, reasons. This sounds almost tautologous, doesn't it? If I try to convince you to perform some action—have dinner with me at a certain restaurant, say—I will, typically, supply you with reasons for doing so; if I try to explain to you why I believe some proposition, I will typically, again, give you my reasons.

In spite of their prevalence in our ordinary talk, however, reasons appear to play no role at all in our two predominant formal theories of reasoning: standard logic, and decision theory. In logic, for example, we learn how to establish that premises imply conclusions, either in accord with formal rules of inference, or through a model-theoretic definition. In decision theory, we learn how to rank proposed actions by expected utility, defined in terms of the possible outcomes of those actions, their probabilities, and utilities. Neither of these well-established theories involves any appeal to our ordinary concept of a reason.

What I try to do in this book is provide a precise, formal account of reasoning that is based on our ordinary notion of a reason—I try to show, in detail, how reasons support the actions or conclusions they do. As you note, I propose that reasons support outcomes in the same way that the premise of a default

rule supports its conclusion. Indeed, I take default logic—or an extended default logic developed for this purpose—as the logic of reason relations, and explore applications in both ethics and epistemology.

Just to give you an example of the issues involved, consider that many people write as if reasons support actions or conclusions only by providing a kind of metaphorical weight. In order to decide whether to pursue one course of action rather than another, for example, it is often thought that what I should do is simply balance the reasons favoring the first course against those favoring the second. But the role of reasons can be much more complicated, as people have noted, but not yet, I don't think, explored carefully. Sometimes reasons can be incomparable in strength, so that what we are really dealing with is a priority ordering, rather than a weighing, among reasons. Sometimes we have to reason about the relative priority to be assigned to other reasons, as when we evaluate evidence about which of two conflicting sources of testimony to take most seriously. And sometimes we have reasons to conclude that another reason is to be given no weight at all, as when a judge tells us that the otherwise persuasive testimony of a witness is inadmissible. The goal of the book is to provide a precise account of how reasons work in cases like these.

## The Unsatisfactory Paradox

The intuition that proper classes are too big to be members, warrants distinguishing them from sets to avoid Russell's paradox. An additional intuition against circularity avoids some other contradictions, including that of Canute's favourite set below. But I will show that there is a similar paradox, the Unsatisfactory paradox, not avoided by these mitigations.

Canute's favourite set is the set of things not in Canute's favourite set (premise). Therefore, the tide is a member of Canute's favourite set iff it is not a member. This follows, given Abstraction, a principle that the members of a set are just those objects of which a predicate is the case. The principle warrants the following schema under the restriction that the open sentence, P$x$, must be $z$-free.

(Abstraction schema) $\forall z\ (z \in \{x{:}\mathrm{P}x\}$ iff $\mathrm{P}z)$

However, there is no reason to think such a set exists for Canute to favour. Even if one exists, there is then no assurance 'Canute's favourite set' has a unique reference, whereas Abstraction will prevent circularity using a free variable. Without a true premise, there is no paradox.

Nevertheless, we could use a function to obtain a true premise. Indeed, we can adapt a self-predication function from Quine (1995: 'Truth, paradox and Gödel's theorem', in his *Selected Logic Papers*, 2nd edn, Harvard University Press, p. 236–41) to yield a set as a result. Here are two examples:

> The self-predicated set of '$y$ is a letter in $x$' = {$y$: $y$ is a letter in '$y$ is a letter in $x$'}.
>
> The self-predicated set of '$y$ refers to $x$' = {$y$: $y$ refers to 'y refers to $x$'}.

The intended function maps a binary predicate to the set of things satisfying the monadic predicate formed by the self-predication of that expression. The predicate can be exhibited and the function represents a purely formal operation; so the function is definable. As an instance of a defined function, premise (1) is *prima facie* true, and defines a constant that can validly be used with Abstraction in the following argument, our Unsatisfactory paradox.

> (1) The self-predicated set of '$y \notin$ the self-predicated set of $x$' = {$y$: $y \notin$ the self-predicated set of '$y \notin$ the self-predicated set of $x$'} by definition of the function.
>
> (2) $\forall z\ (z \in$ {$y$: $y \notin$ the self-predicated set of '$y \notin$ the self-predicated set of $x$'} iff $z \notin$ the self-predicated set of '$y \notin$ the self-predicated set of $x$') Abstraction
>
> (3) $\forall z\ (z \in$ the self-predicated set of '$y \notin$ the self-predicated set of $x$' iff $z \notin$ the self-predicated set of '$y \notin$ the self-predicated set of $x$') (1), (2) =E

The Unsatisfactory paradox is not constrained by the restriction on classes, for it is paradoxical whether anything is a member of the Unsatisfactory class. Nor is the paradox avoided by the free variable constraint on valid use of the Abstraction schema. One might suspect the Law of Excluded Middle as a single cause of paradox, arguing that although (3) has the form A iff ¬A, LEM is required to derive A ∧ ¬A. However, LEM is not explicitly required. Having derived A iff ¬A, assume A, from which ¬A. So, by *reductio* ¬A. From which also A.

An iterative conception of sets avoids these naive paradoxes by insisting sets are formed hierarchically from given objects, not generated by means of other functions. In some ways this seems common sense, but in my opinion it is analogous with the Pythagorean insistence that only those numbers formed by ratios are true numbers, which was then common sense, despite the paradoxical proofs that some square roots were not rationals.

The Unsatisfactory paradox is the set-theoretic analogue of my Unsatisfied Paradox (*The Reasoner* 6(12), 2012, p. 184–5). My favourite predicate just happens to be 'does not satisfy my favourite predicate'. By a Satisfaction schema analogous to Abstraction, it follows that an object satisfies my favourite predicate iff it does not. I argued as follows (*ibidem*). If two paradoxes have one cause and formally analogous reasoning, they have comparable paradoxical extensions. (This is a plausible premise - why would they have radically different extensions?). The Liar, Grelling's paradox and the Unsatisfied paradox seem to have comparable reasoning but their predicates have radically different extensions. Therefore, they have more than one cause or causal factor. I would make the same argument with reference to Russell's paradox and my Unsatisfactory paradox.

Laureano Luna and Martin Cooke have responded with different uniform diagnoses and solutions; they cannot both be right. Luna in 'Satisfiable and unsatisfiable paradoxes: how closely related' (*The Reasoner* 7(5), 2013, p. 56–7) argues circular chains of reference are the sole cause, and should be banned: such expressions are meaningless.

I was able to exhibit my favourite predicate and it is syntactically well-formed. This is sufficient for me to refer to it. (Referring to properties is more challenging, like trying to refer to Canute's favourite set.) So the term 'my favourite predicate' refers to 'does not satisfy my favourite predicate'. This is simply the interpretation of the singular term 'my favourite predicate', not the interpretation of the predicate itself. (Depending on the theory I am responding to, I could use various terms, such as 'what, if anything, is expressed by my favourite

syntactic predicate or predicate symbol'.) In some such standard grammatical way, the reference of a singular term like 'my favourite predicate' is assured. I did not use the Satisfaction schema to define my favourite predicate; it is a contingent fact which predicate happens to be my favourite. Given compositionality, interpreting my favourite predicate itself is then a matter of interpreting 'satisfies'. Luna maintains that my favourite predicate cannot be in the domain of 'satisfies' because my favourite predicate uses 'satisfies'. How then are we to express 'Nothing satisfies my favourite predicate'? If my favourite predicate is not in the domain of 'satisfies', then this is a truth no longer expressible under Luna's constraint, for its assertion would use my favourite predicate, which would mean it has some meaning.

Cooke in 'On the Cause of the Unsatisfied Paradox' (*The Reasoner* 7(6), 2013, p. 69) argues that satisfaction is not always categorical, in effect that $x$ satisfies a predicate, P, insofar as P$x$; and that $x$ satisfies my favourite predicate as much as not. Nevertheless, this uses a different or extended satisfaction relation, and the paradox remains for the categorical relation. This is perhaps clearer for set membership. If Abstraction were re-expressed with 'in so far as' in place of 'iff', the new schema would use a different membership relation or one for which the categorical relation was a limit case. So the Unsatisfactory paradox would still remain paradoxical for the categorical relation.

Peter Eldridge-Smith
Philosophy, ANU

## Defeat reconsidered and repaired

Philosophers commonly distinguish between demonstrative arguments, which are designed to preserve truth, and non-demonstrative arguments, which conserve something else. R.A. Fisher (1936: "Uncertain Inference", *Proceedings of the American Academy of Arts and Sciences*, 71, pp. 245–58), for example, thought that statistical reduction was a type of logical, non-demonstrative inference whose aim is to assign a statistical probability to an individual. What is preserved by Fisher's early proposal for *direct inference* is a presumption of "representativeness" of a statistical class to its individuals. John McCarthy and Pat Hayes (1969: "Some Philosophical Problems from the Standpoint of Artificial Intelligence", *Machine Intelligence*, 4, p. 463–502) thought that non-demonstrative reasoning was necessary in order for a robot to keep track of the things that remain unchanged by its actions. The consequences wrought by opening a door, for instance, do not typically include the shutting of a window. What is preserved by McCarthy and Hayes' non-demonstrative *frame conditions* are commonsense assumptions governing office building kinematics.

John Pollock (1987: "Defeasible Reasoning", *Cognitive Science*, 11, pp. 481–518) thought that a common feature of these and other examples of non-demonstrative reasoning is the provisional standing of their conclusions, and how the evidential support for those conclusions may be *defeated* by additional information (1987: p. 484):

> DEFEATER: Where $D$ and $E$ are jointly consistent propositions, $D$ is a *defeater* for $E$'s support for $H$ if and only if $E$ is a reason to believe $H$ and $E$ & $D$ are not a reason to believe $H$.

In a recent paper, Jake Chandler (2013: "Defeat Reconsidered", *Analysis*, 73(1), pp. 49–51) identifies an unwarranted symmetry constraint in Pollock's definition, namely

> SYMMETRY: For propositions $E$, $D$, and $H$, if both $D$ and $E$ provide a reason to believe $H$, then $D$ is a defeater for $E$'s support for $H$ if and only if $E$ is a defeater for $D$'s support for $H$ (Chandler 2013: p. 50).

Chandler argues, convincingly, that Symmetry should not be a necessary condition for evidential defeat: it is straightforwardly possible for $D$ to defeat the support that $E$ gives to $H$ without $E$ defeating $D$'s support for $H$. Chandler then proposes an alternative to Pollock's account, one that avoids Symmetry. But Chandler's fix, which he calls Defeater′, runs into a difficulty of its own.

> DEFEATER′: Where $D$ and $E$ are jointly consistent propositions, $D$ is a *defeater* for $E$'s support for $H$ if and only if $D$ is a reason to not believe that $E$ is a reason to believe $H$ (Chandler 2013: p. 50).

The problem with Defeater′ is that it fabricates phantom support for a defeater to defeat: $D$ may be a reason to not believe that $E$ is a reason to believe $H$—which thereby suffices for $D$ to defeat $E$'s support for $H$—without $E$ being a reason for $H$ in the first place. For example, inspecting the first 10 light bulbs from a production line and finding all 10 defective ($E$) does not provide a reason to believe that the next bulb off the line is faultless ($H$). Even so, learning that the last delivery of filaments to the factory are all oxidized ($D$) is a reason to believe that the 10 defective light bulbs do not provide a reason to believe that the next bulb is faultless. By Defeater′, $D$ is a defeater for $E$'s support for $H$ even though $E$ does not support $H$.

Defeater′ can be repaired by stipulating, as Pollock does in his original analysis, that $D$ is a defeater for $E$'s support for $H$ only if $E$ gives support for $H$. This yields the following repair to Chandler's Defeater′, namely

> DEFEATER″: Where $D$ and $E$ are jointly consistent propositions, $D$ is a *defeater* for $E$'s support for $H$ if and only if (i) $E$ is a reason to believe $H$ and (ii) $D$ is a reason to not believe that $E$ is a reason to believe $H$.

Defeater″ enjoys all the advantages of Chandler's Defeater′, and does so without the spectacle of phantom support.

Gregory Wheeler
Munich Center for Mathematical Philosophy, Munich
Max Planck Institute for Human Development, Berlin

# News

## Bayesian Blogs

I have recently launched a blog—BayesianWatch—dedicated to the theory and practice of the Bayesian analysis of argument. The theory side will emphasize what Bayesian theory brings to understanding the normative force of arguments (you can see my "Bayesian Informal Logic" in *Informal Logic* (2004) for some of this), including the value of using Bayesian networks for mapping arguments. The practical side will apply argumentation theory to analysing and criticizing natural language arguments as they are found in real life, emphasizing my

own particular interests and biases (e.g., politics and scientific method).

I also run a related blog, Bayesian without Borders. BWB focuses on Bayesian techniques, in particular, the technology and application of Bayesian nets and general issues about Bayesian reasoning, without so much politics or argumentation theory.

Comments and possible contributions are welcome.

Kevin Korb
IT, Monash

## Modeling epistemic and scientific groups: interdisciplinary perspectives, 25–26 November

An interdisciplinary workshop devoted to the study of epistemic groups, be they from everyday life (e.g., friends, family), the economic world (e.g., firms) or science (e.g., research labs, collaborations) was held in Nancy (France) at the MSH Lorraine. It was a sequel of a similar workshop held in 2012 in the same place (cf. *The Reasoner*, 7(2), p. 20), with this time a stronger emphasis on the different modeling and simulation approaches that are currently used to understand epistemic groups. The workshop gathered scholars from various fields such as philosophy of science, epistemology, sociology of science and economics, which fostered rich discussions and exchanges fed by different modelling perspectives.

J. McKenzie Alexander (LSE) argued that the implementation chosen by Weisberg and Muldoon (2009) in their paper about epistemic landscapes in science is not coherent with their hypotheses, and that their results are not robust under the corresponding change in implementation. R. Hegselmann (Bayreuth) presented extensions of the Hegselmann-Krause model, with in particular the addition of network structures to describe exchanges in communities. K. Vaesen (Eindhoven) discussed how models from cultural evolutionary theory could be applied to science and its collaborative features. F. Bloch (École Polytechnique, Paris) presented various recent works in economics on the modeling of groups and coalitions. T. Boyer-Kassem (Lille and Nancy) and C. Imbert (Nancy) presented a model for academic collaboration and its explanation based on deriving groups efficiency from various community structures.

I. Douven (Groningen) argued that within a social setting, inference to the best explanation could be more efficient than Bayesian reasoning. D. Chavalarias (EHESS, Paris) presented simulations of the dynamics of theory inspired by Popper's epistemology, with agents choosing to work on new theories or trying to refute established ones. A. Scharnhorst (The Hague) described various recent mappings of science, and discussed their possible uses. C. Paternotte (LMU, Munich) discussed the virtues used in theory choice, and argued that they could be best understood as catalysts, improving success in the presence of the right ingredients. Finally, C. Proietti (Lund) analyzed the mechanisms of collective epistemic failure, such as pluralistic ignorance, in a simulation framework inspired by works in epistemic logic.

The organizers of the workshop were Thomas Boyer-Kassem (Lille and Nancy), Henri Galinon (Clermont-Ferrand) and Cyrille Imbert (Nancy).

Thomas Boyer-Kassem
Philosophy, Lille and Nancy
Henri Galinon
Philosophy, Clermont-Ferrand
Cyrille Imbert
Philosophy, Nancy

## False-Belief Tasks, 5 December

On December 5, 2013, a workshop on false-belief tasks was held at the University of Amsterdam's (UvA) Institute for Logic, Language and Computation (ILLC). The organizer and host Dr. Jakub Szymanik managed to bring together a group of researchers from different backgrounds, ranging from logic to computer science and cognitive neuroscience.

False-belief tasks are a range of developmental psychological tasks that test whether children are capable of attributing false beliefs to other agents. This ability is a key-ingredient in the child's 'Theory of Mind' (ToM), i.e., the understanding of other agents' mental states. An often-used version of the false-belief tasks is the 'Sally-Anne'-task. During this task, the child listens to a story about a girl named Sally who places her marble in a basket. When Sally leaves the room, Anne takes the marble and hides it in a box. Now Sally returns to her room and the child is asked "*where will Sally look for her marble?*". Children below the age of 4 tend to respond with the factual state of the world (the box) instead of the state of the world as falsely believed by Sally (the basket).

Rineke Verbrugge, Professor of Logic and Cognition at the University of Groningen (RUG), kick-started the topic during another workshop earlier that week on the social dynamics of information change. Verbrugge presented her work on scrub jays, a specific kind of crow, who hide caches of food and occasionally steal from rival caches. As food is an essential pre-requisite for survival, prevention of thievery is tied to Darwinian fitness. Birds protect themselves against theft by employing several strategies; one of them is to re-locate food when it was cached while another bird was watching. Although this kind of behavior seems to imply a ToM, Verbrugge showed that the same behavior can be generated by a virtual bird that operates on behavioristic rules that react to stress induced by the presence of others rather than actual representation of other agents' beliefs (van der Vaart, Verbrugge, & Hemelrijk, 2012: Corvid re-caching without 'Theory of mind': A model. *PloS one*, 7(3)).

The first speaker of the current workshop, Burcu Arslan, a PhD candidate specialising in higher order social cognition of children, affiliated with the artificial intelligence group of the RUG, proposed an ACT-R model of second order false belief reasoning. Second order false belief tasks test whether children can reason about the false beliefs of an agent about the false beliefs of another agent. The learning process of the ACT-R model from first to second order reasoning mimics the developmental progress in children. Future plans were mentioned to test the model's predictions with 5–6 yr olds (Arslan, Taatgen, & Verbrugge, 2013: Modeling Developmental Transitions in Reasoning about False Beliefs of Others. In *Proc. of the 12th International Conference on Cognitive Modelling*).

Thomas Bölander, associate professor from the Technical University of Denmark (DTU), from the department of applied

mathematics and computer science, gave a formalization of the Sally-Anne task in Dynamic Epistemic Logic in order to allow artificial intelligence to be capable of reasoning about other peoples' beliefs and knowledge as this is an essential prerequisite for the proper functioning of socially sensitive service robots.

Stefan Wierda, post-doctoral researcher at the (RUG) specialized in cognitive neuroscience of visual attention, presented another ACT-R model and proposed the possibility that there could be transfer of skills that are engaged in similar yet different tasks. Wierda proposed an analogy between the Sally-Anne task and the Marble Drop Game in which subjects have to engage in higher order reasoning about the goals of other agents instead of false beliefs. Stefan proposed an underlying capacity for inductive reasoning in which the same basic inference rule is re-used several times. An interesting question that remains open is why humans are limited to roughly 3 orders of reasoning.

Torben Braüner, associate professor of computer science at Roskilde University, specialized in logic and its application in computer science, proposed a similar common cognitive basis for different tasks. Braüner noticed a similarity in the logic behind the Sally-Anne task and the Smarties task. The Smarties task shows that children below the age of 4 have not only difficulties with *attributing* false beliefs to *others* but also with *remembering* their *own* previous false beliefs. Braüner formalized this similarity in the form of a natural deduction system for hybrid modal logic in which satisfaction operators can effect jumps between both different time-points and/or agents (Braüner, 2013: Hybrid-Logical Reasoning in False-Belief Tasks. In *Proceedings of Fourteenth conference on Theoretical Aspects of Rationality and Knowledge (TARK)*).

The workshop was concluded by Nina Gierasimczuk, a post-doc researcher at the ILLC, who gave an unannounced surprise talk in which she summarized the workshop and emphasized the importance of interdisciplinary collaboration. As an example of such collaboration, Gierasimczuk presented her work on reasoning strategies of children playing the Mastermind game in which she collaborated with the psychology department of the UvA. She obtained her data from a massive online educational system in which thousands of school children learn arithmetic and abstract reasoning through online games (Gierasimczuk, van der Maas, & Raijmakers, 2013: An analytic tableaux model for Deductive Mastermind empirically tested with a massively used online learning system. *Journal of Logic, Language and Information*).

In conclusion it can be said that the presentations and the ensuing discussions were a success and a nice example of the interdisciplinary cross-pollination that is becoming more and more characteristic of the diverse field of cognitive science.

Gert-Jan Munneke
ILLC, Amsterdam

## Calls for Papers

Causal Discovery and Inference: special issue of *ACM Transactions on Intelligent Systems and Technology*, deadline 14 March 2014.
Presuppositions: special issue of *Topoi*, deadline 15 May 2014.
Virtues & Arguments: special issue of *Topoi*, deadline 1 September 2014.

# What's Hot in ...

## Logic and Rational Interaction

Imagine a traveller is on a long journey to a remote city. Hungry from the long journey, he looks out for a restaurant close to the station. Lucky as he is, he finds two, both open, one full of people, the other completely empty. Presumably, he will go to the full restaurant—the locals will have their reasons.

This is a typical example of a social proof, a belief based on social rather than direct information. It is the same phenomenon that occurs when buying stocks that are on the rise. That a stock is rising *means* that more people are buying than selling—and these people will presumably have reasons too.

The strategy of social proof depends upon various assumptions. One being that others have more information on the subject matter than I do. But they might run the same tactics, basing their behaviour upon my actions as much as my behaviour is based upon theirs. In this case, the emergent group behaviour might be founded on nothing but a big epistemic feedback loop.

Feedback loops depend upon a supply of information on the behaviour of others—information that is nowadays abundantly provided through social networks and other online tools. Or to say it with a term coined by P. Hansen, V. Hendricks and R. Rendsvig: The bigger an infostorm, the likelier it is to contain some local whirls. Under the title infostorm they published an overview article containing thriving examples of feedback phenomena as well as a classification of the various phenomena that can occur based on different kinds of information flows. Crucially, infostorms are not per se rational nor irrational. They can reiterate *correct* beliefs, helping large groups in their epistemic endeavours. But they can also reinforce *false* beliefs and create for instance bubbles at the financial market.

Our traveller from above might have had a tip from a friend recommending the other restaurant. But then, seeing masses of locals all picking the same restaurant, he decides to go along, rather than trusting his friend at home. A fortiori, doing so increases the social information for the traveller walking behind him, also dragging him to the same restaurant, independent of his prior beliefs. This phenomenon is dubbed an informational cascade: A sequential interaction in which the social evidence becomes so strong that it overwrites any contradicting private information the agents might have.

Last month's column describes a quick cascade in a football stadium, creating a universal false belief about a goal within seconds. There, we described all actors involved by Bayesian updating, comparing probabilities of the different options. But what is the logical *structure* of informational cascades, what is the precise mechanism giving rise to cascades?

This question has lately been addressed by various papers. In their work Logical Models of Informational Cascades, A. Baltag, Z. Christoff, J.U. Hansen and S. Smets use tools from Probabilistic Dynamic Logic to model the emergence of informational cascades. Prior to their work, it was known that

rationality does not prevent cascades. A group of individually perfectly rational agents can nevertheless create an information cascade. The analysis in this paper goes even beyond, showing that even higher order reasoning and common knowledge sometimes cannot prevent the emergence of an informational cascade.

In his Aggregated Beliefs and Informational Cascades, R. Rendsvig presents agents as logical machines, equipped with reaction rules for every possible situation they might encounter. The paper analyses when the automated sequential interaction of such logical machines gives rise to a cascade, reaffirming results obtained by other techniques while having the agents rather than the modeller create the dynamics.

LORIWEB is always happy to publish information on topics relevant to the area of Logic and Rational Interaction—including announcements about new publications and recent or upcoming events. Please submit such news items to Carlo Proietti, our web manager via submit@loriweb.org.

Dominik Klein
TiLPS, Tilburg University

## Uncertain Reasoning

One central principle in uncertain reasoning is the requirement that, in solving a well-posed problem, one should use all and only the information which the description of the problem makes explicitly available. All great uncertain reasoners, from Jacob Bernoulli to the present, have contributed to sharpening our thoughts on this rather elusive principle. In recognition of Jacob's pioneering role, Keynes in his 1921 *Theory of Probability* refers to it as *Bernoulli's maxim*. Carnap dubbed it the *Principle of total evidence* whereas Paris named it the *Watt's Assumption*, after the author of an essay eloquently titled "This is IT".

Obvious at it may seem, this principle has often been violated, thereby making room for consequences ranging from the seemingly paradoxical, to the plain wrong. Two well-known examples which sit comfortably in this spectrum are Bertrand's chord problem and Goodman's "Grue".

According to economist Daniel Sgroi, there is one piece of information which should enter the (not-so-well-posed) problem of evaluating the quality of academic research, namely journal-related metrics. His view is developed in detail in the paper D. Sgroi and A. Oswald (2013: "How Should Peer Review Panels Behave?", *Economic Journal* 123(570): 255–278) and summarised in a short piece eloquently titled A Bayesian approach to the REF: finding the right data on journal articles and citations to inform decision-making. Here's the key point, which refers to the case-study of their investigation, namely the UK 'Research Excellence Framework' (REF):

> To cut to the punch-line, experts on REF panels can take a weighted average of the journal's impact factor (representing 'prior information') and the accrued citations to the article (representing 'news'). As the years pass, more and more weight can be placed on the accrued citations. In the early years, almost all the weight can be placed on the journal's impact factor. Bayesian methods will tell us how to adjust the weights as time passes. There is plenty of room for expert opinion and sensible judgement to play a role, so this is not purely mechanistic.

I invite readers with a view on Bayesian updating to have a look at the Sgroi and Oswald (2013) paper for full details. One cannot help noting that, put in this way, the line of argument seems pretty convincing. After all, academics do make choices when it comes to submitting their work for publication, and those choices may well be the output of a process aimed at placing the paper in the 'best possible journal'. And in doing that, people often feel that they know what the 'best journals' are in their own area.

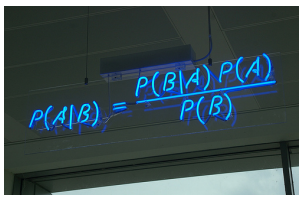


However, there might be a number of concerns related to the Sgroi-Oswald proposal. Perhaps the most obvious one is that both journal metrics and citations are extremely sensitive to manipulation. This is at odds with the requirement that Sgroi points out to the effect that "The number we use for that prior needs to be built from as objective a measure of journal quality as we can find". Although an important part of their research is devoted to "finding the right data", the manipulability of the data itself is likely to be a hard issue to address. In fairness to Sgroi, he concludes the blog piece by pointing out that Bayesian methods should be used as an aid, rather than a substitute for good judgment, in the Peer Review evaluation of academic research. Yet we very often fall for the charm of objective measures, even when all provisos are clearly stated.

This piece is very interesting, on an armchair sociology level, for the way it portrays Bayesian updating:

> This is how economists believe that economic agents merge any prior information they may have with any data they can acquire when trying to solve complex problems in an uncertain world.

As an outsider to economics, my feeling is that economists, especially decision theorists, aren't so lined-up as Sgroi seems to suggest, in defending the ideas we learned to trace back to Reverend Bayes. In addition, some statisticians suggest that even those who label themselves as Bayesians very often fail to practice what they preach (see my November 2011 column for some pointers).

Finally, a friendly escalation towards the frivolous. The short piece was published originally on Vox, under a title which doesn't mention the B-word and with no flashing representation of how beliefs should be updated. I think that's sociologically interesting too.

Hykel Hosni
Marie Curie Fellow,
CPNSS, London School of Economics

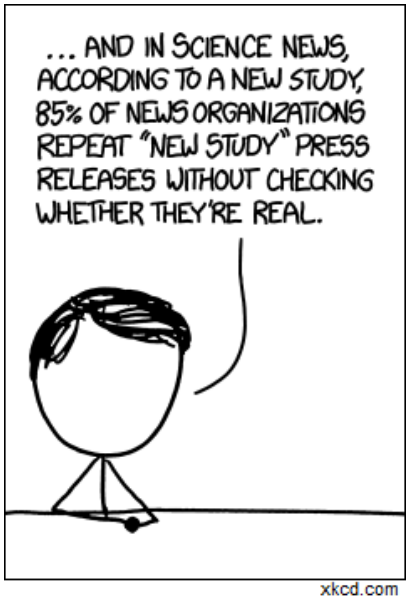


| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | | 8 | | | 5 | | | 4 |
| 1 | | 3 | | 2 | | | 9 | |
| 6 | | 2 | 9 | 4 | | | | |
| 2 | 3 | | 5 | | | | | |
| | | 4 | | 3 | | 9 | | |
| | | | | | 4 | | 5 | 3 |
| | | | | 7 | 8 | 3 | | 9 |
| | 9 | | | 5 | | 6 | | 8 |
| 3 | | | 1 | | | 4 | | |

# Events

## February

PHILOGICA: 3rd Colombian Conference on Logic, Epistemology, and Philosophy of Science, Bogotá, 12–14 February.
Paraconsistency: 5th World Congress on Paraconsistency, Kolkata, India, 13–17 February.
CTFM: Computability Theory and Foundations of Mathematics, Tokyo Institute of Technology, Japan, 17–20 February.
LINZ: Graded logical approaches and their applications, Linz, Austria, 18–22 February.
FPC: Functions, Proofs, Constructions, Tübingen, 21–23 February.

## March

WBEM: Workshop on Beauty and Explanation in Mathematics, Umeøa University, Sweden, 11–12 March.
EMLP: Empirical Methods of Linguistics in Philosophy, Technische Universität Dortmund, 13–14 March.
Justification: Towards an Epistemology of Understanding: Rethinking Justification, Bern, 21–22 March.
CELL: 2nd Logic and Language Conference, London, 21–22 March.
MIND: 9th Mind Network Meeting, Oxford, 26 March.
NNPS: Nordic Network for Philosophy of Science, Lund University, Sweden, 27–28 March.

## April

NAG: Norms, Actions, Games, London, 1–2 April.
AISB: 7th AISB Symposium on Computing and Philosophy: Is computation observer-relative?, Goldsmiths, London, 1–4 April.
HAPOP: History and Philosophy of Programming, Goldsmiths, University of London, 1–4 April.
D& MC: Deductive and Mathematical Cognition, Bristol, 7–8 April.
EBL: 17th Brazilian Logic Conference, Petrópolis, Brazil, 7–11 April.
PSX4: Philosophy of Scientific Experimentation 4, Pittsburgh, PA USA, 11–12 April.
PhiloSTEM: 6th Midwest Workshop in the Philosophy of Science, Technology, Engineering, and Mathematics, Fort Wayne, Indiana, 11–12 April.
Mathematical depth: University of California, Irvine, 11–12 April.
TAMC: 11th Annual Conference on Theory and Applications of Models of Computation, Anna University, Chennai, India, 11–13 April.
LMS: London Mathematical Society Lectures by Jouko Väänänen, London, 14–17 April.
PhML: Philosophy, Mathematics, Linguistics: Aspects of Interaction, St. Petersburg, Russia, 21–25 April.
PhDs in Logic: Utrecht, The Netherlands, 24–25 April.
MAICS: 25th Modern Artificial Intelligence and Cognitive Science Conference, Gonzaga University, Spokane, WA, USA, 26–27 April.
UConnLogic: Abstractionism / Neologicism, University of Connecticut, 26–27 April.
UK-CIM: UK Causal Inference Meeting (UK-CIM): Causal Inference in Health and Social Sciences, University of Cambridge, Cambridge, 28–29 April.
GIRLS: 3rd Conference on Games, Interaction, Reasoning, Learning & Semantics: Evolution and Cooperation, Lund, 28–30 April.
RSC: Research Students' Conference in Probability and Statistics, Nottingham, 28 April–1 May.

## May

LAMAS: 7th Workshop on Logical Aspects of Multi-Agent Systems, Paris, France, 5–6 May.
MSDM: Workshop on Multi-Agent Sequential Decision Making Under Uncertainty, Paris, France, 5–6 May.

SQUARE: 4th World Congress on the Square of Opposition, Pontifical Lateran University, Vatican, 5–9 May.
MS6: Models and Simulations 6, University of Notre Dame, 9–11 May.
EIDYN: Normativity and Modality, Edinburgh, 9–11 May.
Formal Methods: Singapore, 12–16 May.
WPI: 6th Workshop in the Philosophy of Information, Duke University, 15–16 May.
SLACCR: St. Louis Annual Conference on Reasons and Rationality, St. Louis, MO, 18–20 May.
Science & Metaphysics: Ghent, Belgium, 20–21 May.
FilMat: 1st International Conference of the Italian Network for the Philosophy of Mathematics, Milan, 29–31 May.
Formal Ethics: EIPE, Erasmus University Rotterdam, 30–31 May.

June

MSLP: Mathematising Science, University of East Anglia, 1–3 June.
F& MI: Fundamentality and Metaphysical Infinitism, University of Helsinki, Finland, 2–3 June.
ALGMATHLOG: Algebra and Mathematical Logic: Theory and Applications, Kazan, 2–6 June.
CWAP: Normativity of Meaning, Belief and Knowledge, Krakow, Poland, 4–6 June.
LogicMathPhysics: Ontario, Canada, 5–6 June.
POP: 4th LSE Graduate Conference in Philosophy of Probability, London, 6–7 June.
LG& M: Logic, Grammar, and Meaning, University of East Anglia, 7–9 June.
EC: 15th ACM Conference on Economics and Computation, Stanford University, CA, USA, 8–12 June.
Paraconsistency: Paraconsistent Reasoning in Science and Mathematics, Munich, Germany, 11–13 June.
IYSM: International Young Statistician Meeting, Universitá di Cagliari, Italy, 13–14 June.
LOGICA: Hejnice, Czech Republic, 16–20 June.
SILFS: International Conference of the Italian Society for Logic and Philosophy of Sciences, University of Rome "Roma TRE", 18–20 June.
AMSTA: 8th International KES Conference on Agents and Multi-agent Systems—Technologies & Applications, Crete, Greece, 18–20 June.
FEW: 11th Annual Formal Epistemology Workshop, University of Southern California, Los Angeles, CA, 20–22 June.
SEP: 42nd Annual Meeting of the Society for Exact Philosophy, California Institute of Technology, Pasadena, CA, 22–24 June.
CiE: Computability in Europe, Budapest, Hungary, 23–27 June.
SPS: Metaphysics of Science, Lille, 25–27 June.
A & N: The "Artificial" and the "Natural" in the Life Sciences, University of Exeter, 25–27 June.
SPE: Semantics and Philosophy in Europe, Berlin, 26–28 June.
& HPS: Integrated History and Philosophy of Science, Vienna, Austria, 26–29 June.
EGEC: 4th Annual Edinburgh Graduate Epistemology Conference, University of Edinburgh, 27–28 June.
IPSP: Imprecise Probabilities in Statistics and Philosophy, LMU Munich, 27–28 June.
EEN: European Epistemology Network Meeting, Madrid, 30 June–2 July.

July

IACAP: Annual Meeting of the International Association for Computing and Philosophy, Thessaloniki, Greece, 2–4 July.
WCT: workshop on Computability Theory, Prague, 3–4 July.
CICM: Intelligent Computer Mathematics, University of Coimbra, Portugal, 7–11 July.
TiLXIV: Trends in Logic, Ghent University, Belgium, 8–11 July.
FLoC: 6th Federated Logic Conference, Vienna, 9–24 July.
BSPS: British Scoiety for the Philosophy of Science, University of Cambridge, 10–11 July.
DEON: 12th International Conference on Deontic Logic and Normative Systems, Ghent, Belgium, 12–15 July.
SAT: 17th International Conference on Theory and Applications of Satisfiability Testing, Vienna, Austria, 14–17 July.
IPMU: 15th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Montpellier, France, 15–19 July.
LATD: Logic, Algebra, and Truth Degrees, Vienna, 16–19 July.
NMR: 15th International Workshop on Non-Monotonic Reasoning, Vienna, Austria, 17–19 July.
IJCAR: 7th International Joint Conference on Automated Reasoning, Vienna, Austria, 19–22 July.
AUAI: Uncertainty in Artificial Intelligence Conference, Quebec, Canada, 23–27 July.
KRC: Reasoning Conference, Konstanz, Germany, 24–27 July.
LOFT: Eleventh Conference on Logic and the Foundations of Game and Decision Theory, University of Bergen, Norway, 27–30 July.
UCM: Uncertainty in Computer Models 2014, University of Sheffield, 28–30 July.

# Courses and Programmes

## Courses

MLSS: Machine Learning Summer School, Reykjavik, Iceland, 25 April–4 May.
Epistemic Game Theory: EPICENTER, Maastricht University, 12–23 May.
SIPTA: 6th SIPTA School on Imprecise Probabilities, Montpellier, France, 21–25 July.
MCMP: MCMP Summer School on Mathematical Philosophy for Female Students, Munich, Germany, 27 July–2 August.

## Programmes

APhil: MA/PhD in Analytic Philosophy, University of Barcelona.
Master Programme: MA in Pure and Applied Logic, University of Barcelona.
Doctoral Programme in Philosophy: Language, Mind and Practice, Department of Philosophy, University of Zurich, Switzerland.
HPSM: MA in the History and Philosophy of Science and Medicine, Durham University.
Master Programme: in Statistics, University College Dublin.
LoPhiSC: Master in Logic, Philosophy of Science & Epistemology, Pantheon-Sorbonne University (Paris 1) and Paris-Sorbonne University (Paris 4).

Master Programme: in Artificial Intelligence, Radboud University Nijmegen, the Netherlands.
Master Programme: Philosophy and Economics, Institute of Philosophy, University of Bayreuth.
MA in Cognitive Science: School of Politics, International Studies and Philosophy, Queen's University Belfast.
MA in Logic and the Philosophy of Mathematics: Department of Philosophy, University of Bristol.
MA Programmes: in Philosophy of Science, University of Leeds.
MA in Logic and Philosophy of Science: Faculty of Philosophy, Philosophy of Science and Study of Religion, LMU Munich.
MA in Logic and Theory of Science: Department of Logic of the Eotvos Lorand University, Budapest, Hungary.
MA in Metaphysics, Language, and Mind: Department of Philosophy, University of Liverpool.
MA in Mind, Brain and Learning: Westminster Institute of Education, Oxford Brookes University.
MA in Philosophy: by research, Tilburg University.
MA in Philosophy of Biological and Cognitive Sciences: Department of Philosophy, University of Bristol.
MA in Rhetoric: School of Journalism, Media and Communication, University of Central Lancashire.
MA programmes: in Philosophy of Language and Linguistics, and Philosophy of Mind and Psychology, University of Birmingham.
MRes in Cognitive Science and Humanities: Language, Communication and Organization: Institute for Logic, Cognition, Language, and Information, University of the Basque Country, Donostia, San Sebastián.
MRes in Methods and Practices of Philosophical Research: Northern Institute of Philosophy, University of Aberdeen.
MSc in Applied Statistics: Department of Economics, Mathematics and Statistics, Birkbeck, University of London.
MSc in Applied Statistics and Datamining: School of Mathematics and Statistics, University of St Andrews.
MSc in Artificial Intelligence: Faculty of Engineering, University of Leeds.

MA in Reasoning

A programme at the University of Kent, Canterbury, UK. Gain the philosophical background required for a PhD in this area. Optional modules available from Psychology, Computing, Statistics, Social Policy, Law, Biosciences and History.

MSc in Cognitive & Decision Sciences: Psychology, University College London.
MSc in Cognitive Science: University of Osnabrück, Germany.
MSc in Cognitive Psychology/Neuropsychology: School of Psychology, University of Kent.
MSc in Logic: Institute for Logic, Language and Computation, University of Amsterdam.
MSc in Mind, Language & Embodied Cognition: School of Philosophy, Psychology and Language Sciences, University of Edinburgh.
MSc in Philosophy of Science, Technology and Society: University of Twente, The Netherlands.
MRes in Cognitive Science and Humanities: Language, Communication and Organization: Institute for Logic, Cognition, Language, and Information, University of the Basque Country (Donostia San Sebastián).
Open Mind: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.
PhD School: in Statistics, Padua University.

# Jobs and Studentships

## Jobs

Post-doc Position: in Set Theory, Torino University, until filled.
Post-doc Position: on the project "Rational reasoning with conditionals and probabilities", MCMP, LMU Munich, until filled.
Post-doc Position: on the project "Grounding—Metaphysics, Science, and Logic, Hamburg, deadline 5 February.
Post-doc Position: in Mathematical Logic, Mathematics, Padua, deadline 10 February.
Lecturer: in Theory of Science, University of Gothenburg, Sweden, deadline 13 February.
Post-doc Positions: on the project "The Causal Power of Information in a Quantum World," Philosophy, University of Queensland, deadline 19 February.
Assistant Professor: in Theoretical Philosophy, Salzburg, Austria, deadline 26 February.
Professor: in Analytic Philosophy, Vienna, deadline 31 March.

## Studentships

Student Assistant: on the project "Rational reasoning with conditionals and probabilities", MCMP, LMU Munich, until filled.
PhD Position: in Philosophy of Science, University of Edinburgh, deadline 3 February.
PhD Position: and *The Reasoner* Editorial Assistant, Philosophy, Kent, deadline 3 February.
PhD Position: in Philosophy of Science or Logic, Salzburg, Austria, deadline 12 February.
PhD Positions: at the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems (AIMS), Oxford, deadline 14 February.