THE REASONER

VOLUME 6, NUMBER 12
DECEMBER 2012

CONTENTS

EDITORIAL

I had the privilege to meet Professor Atocha Aliseda

when I was a graduate student at the National Autonomous University of Mexico (UNAM), where she was at that time coordinator of the Philosophy of Science Postgraduate Program. She was very busy then, but was nonetheless kind enough as to listen to newcomers with strange concerns. She leans toward studying specific kinds of applied reasoning, and I'm still looking for very general theories of logic regardless of their uses to reason in the real world, but we both are interested in logic and it is always a pleasure to find time to talk about it.

Atocha Aliseda is full time professor at the Institute for Philosophical Research, National Autonomous University of Mexico (UNAM) and has been visiting professor at the universities of Seville, Salamanca (Spain), San Sebastian (Basque Country), Groningen (The Netherlands), and several other universities in Mexico. She is author of *Abductive Reasoning: Logical Investigations into the Processes of Discovery and Explanation*, as well as of numerous papers on abductive reasoning, logics of scientific discovery, and the philosophy of Charles Sanders Peirce.

I am delighted to be guest editor of this issue of *The Reasoner* and open it with an interview with a Mexican logician. I hope you will enjoy this interview as much as I did.

LUIS ESTRADA-GONZÁLEZ
Philosophy, University of Tartu

FEATURES

Interview with Atocha Aliseda

Luis Estrada-González (LEG): Could you tell our readers about your intellectual history?

Atocha Aliseda (AAL): I hold a bachelor degree in mathematics from UNAM and a PhD in philosophy and symbolic systems from Stanford. During my undergraduate studies I was interested in computer science, but at that time (mid-Eighties) bachelor programs in that area were nonexistent in Mexico. People with interests like mine studied mathematics, where it was possible to get a heavy load of logic courses and some computing ones. My bachelor thesis was a computer program (in Prolog) for the teaching of propositional logic. Soon after I graduated, I became research assistant to a philosopher, and wrote with him my first paper on automated reasoning. Before that, my first job at the university was as editorial assistant in *Mathesis*, a journal in the history and philosophy of mathematics.

As you might have already noticed, I was (and still am) interested in many subjects, so it was not clear what I should do after my undergraduate studies. I was of course interested in mathematics, but was not really aiming at a doctorate in mathematics. I was interested in philosophy and in linguistics. I

did not know I was indeed looking for an interdisciplinary program, until I found (through the post and libraries like the one at the American Embassy in Mexico, in a time with no internet) *Philosophy and Symbolic Systems* at Stanford and *Pure and Applied Logic* at Carnegie Mellon. Symbolic Systems at Stanford turned out to be a great option for me. It is run by the Philosophy Department, but it involves others as well (Linguistics, Psychology, Mathematics, and Computer Science).

I finished my doctorate and returned to Mexico in 1997 as a researcher at the UNAM's Institute for Philosophical Research. Soon thereafter, I attended the Congress of Logic, Methodology and Philosophy of Science held in Cracow in 1999 and met Theo Kuipers there. He knew my thesis via Johan van Benthem, my supervisor at Stanford, and I told him I was interested in continuing my work on abduction, especially the connection to issues in the philosophy of science. At that time, my home university was closed, due to a political turmoil. Sometime after our first encounter, Theo invited me to fill a one-year postdoc position at Groningen. I was desperate to continue my research and the strike at UNAM seemed endless, but I could not just leave my job behind. I suggested to Theo a somewhat unusual scheme of postdoc stay and, surprisingly, both him and my home university accepted. It turned out to be a three-year collaboration (2000–2002), three months every year at Groningen. I have wonderful memories of my time there.

LEG: Maybe some of our readers know of your work on abduction and how you applied it to some issues in the philosophy of science. We will address philosophy of science later, but given what you said about your intellectual history, could you tell us how you became interested in abduction?

AAL: I thank you for this question, for in general one cannot guess how it is that a researcher became interested on a topic merely by looking at her publications.

I attended several logic courses given by van Benthem at Stanford and they were awesome. I discovered then there were many logics beyond classical logics, but my interest in abduction came from linguistics. Following a suggestion by Tom Wasow, I read "Interpretation as abduction", by Jerry Hobbs and others. They claimed that the interpretation of a text can be accounted for by abduction: it is the minimal explanation of why the text would be true. I loved this idea. That was the kind of stuff I wanted to investigate in my thesis, mainly from a logical and a computational perspective.

My first attempt, however, was something impossible: to formalize Charles Peirce's notion of abduction. But his notion of abduction is neither a single nor a clear-cut one, for his thoughts on the subject evolved throughout his life, and his reflection on abduction is so closely tied to the rest of his epistemology and semiotics, that it seems impossible to formalize it. I then decided to take Peirce rather as an inspiration, and put forward a characterization of my own.

LEG: What was your PhD thesis about exactly?

AAL: In my dissertation "Seeking Explanations: Abduction in Logic, Philosophy of Science and Artificial Intelligence", maybe one of the first PhD theses about abduction, my aim is to lay down logical and computational foundations in order to explore some of the formal properties under which abductive logics may be generated and evaluated.

This approach naturally led me to connections with theories of explanation in the philosophy of science and to computationally oriented theories of belief change in artificial intelligence.

LEG: I understand you applied your research on abduction to the philosophy of science, especially to the processes of discovery and explanation. What are your views on this, and how have they evolved?

AAL: As a result of my postdoc research, I wanted to write a book on abduction, with some new results on the logical part as well as on the philosophical one.

When I came to work with Theo Kuipers he told me something along the following lines: "Your model is very nice, but it is not good for philosophy of science because it is merely propositional. You need first order logic to represent scientific laws." He was right. There were as well some issues in the literature on logics of discovery in the philosophy of science that I needed to study further. Working at UNAM in the postgraduate program in philosophy of science was very useful to me because I learned a lot by talking to my philosophy colleagues. At the same time, I had an interlocutor in Seville, Ángel Nepomuceno, who had been working on second order logics and tableaux, and had some ideas as to how to modify tableaux in order to accommodate abduction in (some fragment of) first order logic. I visited him a couple of times while I was in The Netherlands, and we started a collaboration on this issue.

Now, about abduction in discovery and explanation: On the one hand, Hans Reichenbach proposed a methodological distinction between the context of discovery and the context of justification, back in the Thirties of last century. This distinction served very well the positivist focus on justification issues from a logical point of view, but completely left out of the picture anything having to do with discovery. Fortunately, there are always philosophers working off-track well-accepted ideas, and in this case, both Lakatos and Hanson raised the importance of heuristics, very much connected to abduction. But going back to the 19th century, there is of course Peirce, who thought of abduction as the logic for "synthetic reasoning", a method for the generation of new ideas, as the adventuresome logic that gives us plausible hypotheses as the output of logical inferences. So, my book extends my dissertation in that I place my analysis on abduction within the discussion of logics of discovery, present my new results in logic, and included a chapter on Peirce and pragmatism, as well as an analysis of abduction connected to progress in empirical sciences, that I specifically worked on with Theo Kuipers.

LEG: What are the issues, problems or questions you are currently working on?

AAL: I have been working recently on the logical aspects of medical diagnosis and, more broadly, in the philosophy of medicine. At some point, I became very interested in exploring whether my logical models work in the real world, and I found a group of neurologists interested in the process of medical diagnosis, as done by doctors in their discussion sessions. So, I have found a place at the National Institute of Neurology and Neurosurgery in Mexico City, in which I observe physicians in action. It is now clear to me that medical reasoning is neither mainly nor solely abductive, but it is rather an interaction of several types of reasoning (analogical, deductive, inductive), and therefore any formal model must take that into account. Moreover, hypothesis generation is not only via abduction. Doctors bring in hypotheses in the analysis of a case, some-

times for the sole purpose to test them. At the moment, I am working with a computer science master student in using adaptive logics to model some cases of medical diagnosis. I talked about our preliminary results in the conference Model-Based Reasoning in Science and Technology held in Italy as well as at the Evidence and Causality in the Sciences conference in Canterbury.

LEG: In "Mathematical Reasoning vs Abductive Reasoning: A Structural Approach" you oppose abductive reasoning to mathematical reasoning. Isn't mathematical reasoning, like medical reasoning, several things at once too, including abductive reasoning?

AAL: In that work, I compare the classical notion of consequence to that of (one version of) abductive inference, and I do so within the framework of structural rules. There is a representation theorem stating that monotonicity, contraction, reflexivity and cut characterize classical inference. I proved there is also a representation theorem stating that some other forms of above rules (conditional reflexivity, simultaneous cut and conclusion consistency) characterize what I called "consistent abduction", a form of abduction with an underlying classical notion of consequence plus the consistency requirement. So, it is in this context that I draw the distinction, but of course I think that mathematical reasoning involves many thinking patterns and forms, one of them being abductive, indeed.

LEG: You have been teaching recently about creativity, especially computational creativity. How do you characterize that? How does it fit in with your previous research?

AAL: I am a member of the computational creativity seminar at UNAM. It is a seminar for people from postgraduate programs in computer science, engineering and philosophy of science. "Creative" in this context means that of which there is consensus that it is novel and useful. This is of course a limited and controversial characterization, but it serves as a working hypothesis for our seminar.

The course is divided into two parts: Philosophical and Computational. I am in charge of the first one, and my colleague Rafael Pérez y Pérez, one of the pioneers of the study of computational creativity in Mexico, of the second one. After introducing the discussion on discovery in philosophy and cognitive psychology, the key question leading our discussions in the first part is whether machines can generate new knowledge. In the second part we analyze programs of computational creativity, which can be classified roughly into three kinds: those for creative support, those to simulate creative processes and those aimed at solving problems creatively. It is a very stimulating environment bringing together students from the humanities and the sciences.

LEG: Sadly we are reaching the end of the interview. To close it, are there any particular topics in your areas of interest that you would recommend to graduate students starting out today?

AAL: Besides recommending my own topics of interest (philosophy of medicine, creativity, non-classical logics), the idea I want to get through is to encourage students to engage in interdisciplinary work, but without losing sight of the fact that our academic world is based on certain divisions. Although it may sound great to be "an interdisciplinary researcher", someone able to talk to people in several areas, studying a topic concerning several disciplines, faces some difficulties. The obvious one is that one cannot master all disciplines involved in a topic. Moreover, when it is time to publish, get

a job or find funding, we need to make our research fits into a specific discipline, book series or journal. The case of logicians is a case in point: Even though we may be able to talk with certain degree of confidence to researchers in other areas, perhaps we are neither mathematicians, nor philosophers nor computer scientists in full. So the advice is: Engage in interdisciplinary research, but choose one of those disciplines to be the core of your studies and research.

Reasoning without Contradiction

Adding or subtracting a tautology to its premises will have no effect on the validity of an argument, so it is reasonable to believe that tautologies are not required for reasoning. But contradictions, it seems, feature in tried and trusted proof procedures, so one might suppose that, were contradiction somehow banned from proof procedures, incompleteness would result. But I think that this is untrue and that a case can be made that reasoning without contradiction has some very beneficial consequences. Limitations of space prevent me giving more than a sketch, but I set out the argument in such a way that the reader can easily identify where, if anywhere, s/he thinks it goes wrong:

1. The utterance 'It is raining or it is not' literally gives no information at all about the weather (Wittgenstein, *Tractatus* 4.461) and provides the hearer no guide to action, such as to take or not take an umbrella.

2. The utterance says nothing about anything else. (If it is not about the weather, what is it about?)

3. Therefore it says nothing at all; *aliter*, is completely uninformative.

4. *A fortiori*, it says nothing true and nothing false, and a similar conclusion will hold for any tautology.

5. Therefore its negation says nothing true and nothing false (think of prefixing a meaningless word-salad with 'Not').

6. The negation of a tautology is a contradiction.

7. So equally a contradiction says nothing. In particular,

8. A contradiction says nothing false.

This is Wittgenstein's thesis, enunciated in his dictation to Moore in Norway, 1914 (*Notebooks 1914–16*: 108), in his 1939 Cambridge *Lectures on the Foundations of Mathematics* (p. 222) and in his late manuscripts, where he insists that, although the sentence 'It is raining and it is not' can be used to mean that there is a very light drizzle, there can be no contradictions where one deliberately speaks against (contra-dicere) oneself, uttering a sentence of the form $\ulcorner A \& \neg A \urcorner$ and meaning it literally. Such contradictions have no sense; we have no clear-cut use for them (*Remarks on the Philosophy of Psychology, Vol. 2*: §290).

But surely, it might be said, we make frequent use of contradictions in reasoning. Well, in classical logic, we derive conclusions (in fact, any conclusion) from contradictions. But if a contradiction says nothing, and nothing follows from nothing, then we need to abandon the classical principle *ex contradictione quodlibet* as being not characteristic of good reasoning. Several logics do just that and, as Graham Priest has shown ('Negation as Cancellation, and Connexive Logic', *Topoi* 18 (1999): 141–148), consistent connexivist logics do, and are tied to a view of negation as cancellation that is, in my view, attractive and natural. The rule &-E is valid except where the conjunction contains a wff and its negation.

Reductio ad Absurdum is a healthy rule of inference unaffected by the claim that contradictions say nothing. Informally, the rule is: if you want to get something useful (non-vacuous) from an inconsistent set of premises, eliminate at least one of those premises. The contradiction reached in a *reductio* proof is not *used*; it simply serves as a sign warning us to turn back and revise the argument in the way just suggested. Sextus Empiricus (c. 160–210 AD) expressed this beautifully: 'Suppose there were a road leading up to a chasm, we do not push ourselves into the chasm just because there is a road leading to it but we avoid the road because of the chasm; so, in the same way, if there should be an argument which leads us to a confessedly absurd conclusion, we shall not assent to the absurdity just because of the argument but avoid the argument because of the absurdity' (Sextus Empiricus, *Outlines of Pyrrhonism*, 1933, Vol. 1, p. 252).

Reasoning does not require contradiction.

A corollary:

9. Nothing follows from nothing, in that you can get no information from no information.

10. $\ulcorner((x)(x \in R \& \neg(x \in x))\urcorner$ is a contradiction—instantiate x with R.

11. Therefore nothing follows from it.

12. Therefore in any logic (e.g., classical) in which $\ulcorner A \& \neg A\urcorner$ is truth-functionally, or informationally, equivalent to $\ulcorner A \leftrightarrow \neg A\urcorner$, nothing follows from $\ulcorner((x)(x \in R \leftrightarrow \neg(x \in x))\urcorner$.

13. Russell proved that Frege's *Grundgesetze* system is inconsistent by invoking the impredicative function $\ulcorner\neg(x \in x)\urcorner$. Russell's proof did not establish the existence of a Russell class R, because *nothing*, including the existence of R, follows from $\ulcorner((x)(x \in R \leftrightarrow \neg(x \in x))\urcorner$.

14. In general, in any definition of a set by the abstraction principle $\ulcorner(\exists y)(x)(x \in y \leftrightarrow Fx)\urcorner$, we must forbid substituends for x, y and F in the scope $\ulcorner(x \in y \leftrightarrow Fx)\urcorner$ such that the result of any such substitution entails a tautology or a contradiction; for tautologies and contradictions say nothing and, in particular, fail to state conditions or give a recipe for determining whether any given object belongs in the set. The Russellian 'definition' $\ulcorner(\exists y)(x)(x \in R \& \neg(x \in x))\urcorner$, as we have seen, falls foul of this prohibition, but there is, for example, nothing to prohibit the definition of the empty set as $\ulcorner(\exists y)(x)(x \in y \leftrightarrow \bot)\urcorner$.

15. Naïve set theory (which surely captures rather faithfully our intuitions about sets), plus some principled, intuitive restriction, is superior to a philosophically ill motivated axiomatic system. The Wittgenstein-inspired proposal here is motivated not by the need to avoid contradiction, but by the requirement to avoid illicit definition.

LAURENCE GOLDSTEIN
Philosophy, Kent

The Unsatisfied Paradox

Just as Russell's, Cantor's and Burali-Forti's paradox pose foundational challenges in set theory, semantic paradoxes pose philosophical challenges for a theory of truth. Two traditional paradoxes of truth are the Liar (4th century BCE), using the one-place truth predicate, and Grelling's (1908) using the truth relation. They are seemingly sound arguments to contradictions. There is still no consensus on what is wrong with their reasoning: various analyses of these have led to many philosophical theories of truth. Nevertheless, paradoxes of truth are widely believed to use the same kind of pathological reasoning and to stem from the same cause, whatever that cause may be (for example, some vicious circularity or attributing truth to sentences that do not express propositions). Three is a crowd though: the Unsatisfied paradox (Eldridge-Smith 2008: *The Liar Paradox and its Relatives*, Australian National University), as another type of paradox of the truth relation, disconfirms the assumption of a common or single cause. I will present several versions of the newer paradox and contrast it with Grelling's, then conclude by outlining the argument against assuming a common cause. This warrants further research in truth theory.

In parallel to Tarski's derivation of the Liar (1935: 'The Concept of Truth in Formalised Languages' in his 1983 *Logic, Semantics, Metamathematics*. Indianapolis: Hackett, 158), let the symbol '*p*' abbreviate '*the first predicate exhibited in this article*'. Consider the following predicate:

 does not satisfy p

Having regard to which predicate contingently is first exhibited in this article:

 (1) *p is identical with* '*does not satisfy p*'.

Satisfaction is a relation between objects and predicates. The relation's principled use suggests elimination and introduction rules for 'satisfies': If a thing satisfies a predicate, then that thing is as the predicate says; and conversely, if a thing is as a predicate says, then that thing satisfies that predicate. The following schema combines these rules, where an angle bracket expression is used to represent a canonical name for an enclosed predicate *P* (cf. McGee, V. 1991: *Truth, Vagueness, and Paradox*. Indianapolis: Hackett, 31):

(Satisfaction schema) $(\forall z)(z$ satisfies $\langle P \rangle$ iff $Pz)$

(A canonical name is just one from which the predicate can be effectively recovered.) So, for the quotation mark name of the predicate p, for any z:

(2) *z satisfies 'does not satisfy p' if and only if z does not satisfy p.*

The premises (1) and (2) give a contradiction by substitution of identicals:

(3) *z satisfies p if and only if z does not satisfy p.*

It never mattered what z was. The above paradox results for any object.

I devised the Unsatisfied paradox in 2003. Informally:

> *My favourite predicate just happens to be 'does not satisfy my favourite predicate'. Crete satisfies 'does not satisfy my favourite predicate' iff Crete does not satisfy my favourite predicate. Therefore, Crete satisfies my favourite predicate iff Crete does not satisfy my favourite predicate.*

It is paradoxical whether Plato, snow or anything satisfies my favourite predicate.

Having an *insatiable* predicate distinguishes an Unsatisfied paradox from versions of Grelling's and other paradoxes, whose predicates are often otherwise well-behaved. For example, a version of Grelling's paradox concerns whether 'does not satisfy itself' satisfies itself. Substituting '*does not satisfy itself*' for '*P*' in the Satisfaction schema yields (as per McGee *ibid.*, 31):

(4) $(\forall z)$ (*z satisfies 'does not satisfy itself' iff z does not satisfy itself*)

then instantiating z with the same expression yields:

(5) '*Does not satisfy itself*' *satisfies* '*does not satisfy itself*' *iff* '*does not satisfy itself*' *does not satisfy itself.*

While this instantiation of z leads to paradox, many expressions such as 'is German' and 'is edible' unproblematically do not satisfy themselves. The predicate of a Grelling's paradox seems problematic just in reflexive cases; it is not *insatiable*.

An object satisfies a predicate just if that predicate is true of that object, so we can also use the following Relational Truth schema:

(RT-schema): $\langle P \rangle$ *is true of z iff Pz*

And the same paradox can be given using true of, as:

> *My second favourite predicate just happens to be 'my second favourite predicate is not true of'.*

It is paradoxical whether my second favourite predicate is true of Aristotle, Athens or any object. Again, contrast this with Grelling's:

'Heterological' was defined as meaning 'not true of self'; we can therefore ask if the adjectival phrase 'not true of self' is true of itself. We find that it is if and only if it is not, hence that it is and it is not (Quine W. V. O. 1976: *The Ways of Paradox*, 2nd edn, Cambridge, Massachusetts: Harvard University Press, 4).

Other techniques that are commonly used to derive versions of the Liar work as well to derive versions of the Unsatisfied paradox. Consider Gödel's lemma:

For any formula $\phi(x, v_1, \ldots, v_n)$ of $\mathcal{L}$, one can find a formula $\psi(v_1, \ldots, v_n)$ so that $R \vdash (\forall v_1), \ldots, (\forall v_n)(\psi(v_1, \ldots, v_n) \leftrightarrow \phi(\overline{\ulcorner \psi \urcorner}, v_1, \ldots, v_n))$ (McGee *ibid*., 24, with my overlining.)

($\mathcal{L}$ is a language and R is a theory of arithmetic.) The corner bracket expression represents a (suitable) numerical encoding of the contained predicate, and overlining represents its numeral. (That is, $\ulcorner P \urcorner$ represents a canonical name for a predicate P, because it represents the numeral denoting the number that encodes the predicate P.) Accordingly, with *x is not true of z* for $\psi(x, z)$, the lemma ensures that there is a formula, Gz, such that:

(6) $\forall z(Gz \leftrightarrow (\overline{\ulcorner G \urcorner}$ *is not true of z*$))$

But, by the RT-schema:

(7) $\forall z(Gz \leftrightarrow (\overline{\ulcorner G \urcorner}$ *is true of z*$))$

(So the language cannot consistently define its own truth relation.) Nevertheless, if both the above seem true—and they do: (6) on the basis of the proof of Gödel's lemma and (7) on the basis of the Relational Truth schema—we have our Unsatisfied paradox. Again, the instantiation of z is arbitrary.

Distinguished by the paradoxical extensions of their predicates, we have two paradoxes of the truth relation. In terms of formal proofs, the Unsatisfied paradox exactly parallels Tarski's proof of the Liar; Grelling's is actually analogous to Russell's paradox (Quine ibid.,11). Semantically, the predicates of the Liar and Grelling's can be partially defined without paradox; yet the predicate of the Unsatisfied paradox cannot be partially defined, even though it seems to make some sense. If these paradoxes had a common cause, formally analogous reasoning would result in analogous paradoxical extensions. So, they seem to have more than one cause or causal factor. Tarski's analysis of the Liar motivated his approach to defining truth, an analysis from which many modern truth theories draw some elements. In particular, Tarski's inductive definition of truth uses the satisfaction relation. Researching the paradoxical reasoning and inter-relationship of the two paradoxes of the truth relation could inform a development in truth theory.

PETER ELDRIDGE-SMITH
Philosophy, ANU

Causality, Probability, and Time, by Samantha Kleinberg, Cambridge University Press, 2012

Causality is a key part of many fields and facets of life, from finding the relationship between diet and disease to discovering the reason for a particular stock market crash. Despite centuries of work in philosophy and decades of computational research, automated inference and explanation remain an open problem. In particular, the timing and complexity of relationships have been largely ignored even though this information is critically important for prediction, explanation, and intervention. However, given the growing availability of large observational datasets, including those from electronic health records and social networks, it is a practical necessity.

This book presents a new approach to inference (finding relationships from a set of data) and explanation (assessing why a particular event occurred), addressing both the timing and complexity of relationships. The approach allows one to infer a relationship such as "smoking and asbestos exposure until a particular genetic mutation occurs causes lung cancer with probability 0.6 in between 1 and 3 years" without any prior knowledge of a connection between these variables or its timing. In addition to this type of type-level finding, the book introduces methods for explanation with incomplete and uncertain information, building on type-level knowledge while allowing that individual cases may deviate from this. The practical use of the method developed is illustrated through theoretical and experimental case studies, demonstrating its feasibility and success.

More information on the book and the data used in it are available here.

Samantha Kleinberg
Computer Science, Stevens Institute of Technology

Debate Dynamics: How Controversy Improves Our Beliefs, by Gregor Betz, Springer, Synthese Library 357, 2012

Are controversial debates consensus- and truth-conducive? And if so, why? These are the central questions *Debate Dynamics: How Controversy Improves Our Beliefs* seeks to answer. It employs a bounded rationality model of complex argumentation (where proponents are, in particular, not assumed to be logically omniscient) and multi-agent simulations (see this movie) to assess under which conditions the exchange of arguments fosters mutual agreement and makes proponents track down the truth.

The main findings, the argumentation-theoretic model employed, and potential objections are presented and discussed informally in a comprehensive introductory chapter (which is freely available here). That discussion refers systematically to the 14 more specific chapters which report and analyse the detailed, quantitative results of different simulation experiments.

Some *key findings concerning the consensus-conduciveness* of complex argumentation are: (1) On average and in the long-run, controversial argumentation fosters mutual

agreement. But the medium-term effect and the pace of the overall rapprochement depend sensitively on the argumentation strategies employed by the proponents. (2) It is, in general, more consensus-conducive to argue in an opponent-sensitive way, i.e., to introduce arguments whose premises are agreed upon by other proponents. (3) The effectiveness of critical and aggressive argumentation strategies in generating consensus depends on whether the initial agreement with one's opponent is very high ("friend") or very low ("fundamentalist").

Some *main results apropos the truth-conduciveness* of controversial debate read: (1) In toto, proponents approach the truth in a controversy. (2) It's criticism and pluralism that turn out to be the main drivers of such epistemic progress. (3) The most consensus-conducive strategies are not necessarily the most truth-conducive ones: e.g., the more critical an argumentation, the greater the verisimilitude increase. (4) Broad consensus, stability of a proponent position in the face of severe criticism, and degree of justification are reliable indicators of verisimilitude at an early stage of a debate.

Gregor Betz
Philosophy, Karlsruhe Institute of Technology

Computational Models of Argument, 10–12 September

From 10 to 12 September, the Vienna University of Technology hosted COMMA 2012. After three successful editions in Liverpool (COMMA 2006), Toulouse (COMMA 2008) and Desenzano del Garda (COMMA 2010), the Fourth International Conference on Computational Models of Argument lived up to the high standard set in the years before in every way, attracting 70 participants.

The conference was organized by Bart Verheij (Institute of Artificial Intelligence, University of Groningen), Stefan Szeider and Stefan Woltran (both from Vienna University of Technology). COMMA 2012 was part of the Vienna Logic Weeks, a concentration of academic events from logic-related areas during two weeks in September 2012. As such, it offered the opportunity of a plenary talk by the eminent logician and computer scientist Robert Kowalski, who spoke about his work towards a logic-based, unifying framework for computing.

The focus of COMMA 2012 on computational aspects of argumentation was supported by three invited talks which included Trevor Bench-Capon (University of Liverpool), Erik Krabbe (University of Groningen) and Keith Stenning (Universities of Edinburgh and Giessen). Trevor Bench-Capon presented an overview of his work on argumentation over the past forty years, connecting it to insights he gained from philosophy and his experiences in the Civil Service. He applied this combined perspective to several aspects of computational argumentation, such as knowledge based systems, explanation, context, audiences, schemes and models. Erik Krabbe drew a line from the Aristotelian procedure for academic debates through formal dialectic and pragma-dialectics to computational models of argument, establishing that formal dialectic is rooted in ancient dialectic. Keith Stenning's talk concentrated on a multiple-logics view of cognition, showing that several logics influence peoples' reasoning. He discussed the potential consequences of such a view of cognition on theories of argumentation with

reference to real-word examples for conversation and legal institutionalization.

The multi-faceted program of COMMA 2012 included in total 45 technical presentations, a system demonstration session organised by Adam Wyner (University of Liverpool), and a panel on "The added value of argumentation". COMMA 2012 was complemented by a joint dinner at Vienna's Museum of Natural History which presented an excellent opportunity for networking and information exchange. The proceedings of the conference are available here.

KATARINA JURIK
VCLA, Vienna University of Technology

Formal Ethics, 11–13 October

The international conference Formal Ethics 2012 took place at the Munich Center for Mathematical Philosophy (MCMP/LMU Munich) on October 11th to 13th, 2012. The aim of this conference was to bring together researchers interested in applying formal methods in ethical theory. Twenty speakers presented their work on matters from deontic- and (X)stit-logic, (evolutionary) game theory, social choice, theory of action/reasons, axiology, legal theory and decision theory relevant to ethical issues. Keynote speakers were John Horty, Gerhard Schurz, Martin van Hees, and Simon Huttegger.

Day One started off with John Horty on 'Common Law Reasoning', who investigated the mechanisms and justification of rules in legal theory. He introduced a formal framework which explains how it is possible for rules to constrain decisions although these rules seem to be changeable at will. Horty argued that this is due to early cases which build a priority order on the reasons, which develop over time.

Day Two began with Gerhard Schurz on 'Non-trivial Versions of Hume's Is-Ought Thesis'. According to Prior's Paradox, is-ought inferences seem to be possible, if we allow for mixed statements (e.g., p. Therefore: p or Oq). What gives normative content to those mixed statements is the so called *replacement criterion of inessentiality*. This criterion plays a prominent role in a general formulation of Hume's Thesis, which states that every mixed conclusion which is inferable from descriptive premises is normatively irrelevant in a sense which Schurz among other things specified in his talk.

Martin van Hees presented joint-work with Matthew Braham on 'The Inconsistency of Pure Libertarianism'. In order to answer the question whether libertarianism is a consistent political theory, van Hees presented a formal framework whose central notion is that of a *rights structure*, i.e., a mapping from individuals to rights (sets of subsets of outcomes). For libertarianism to be consistent the rights structure has to satisfy four intuitive conditions: completeness, conclusiveness, non-imposition and symmetry. Van Hees and his colleague prove that there exists no rights structure that satisfies all four conditions, and that dropping either one of the conditions, however, yields equally undesirable consequences. They conclude that there is no attractive strategy to rescue pure libertarianism, and hence that no theory of justice could be exclusively based on rights.

On Day Three Simon Huttegger presented work he and Rory Smead have done on 'The Evolution of Norms in Structured Populations'. Huttegger applies the Maynard

Smith Haystack Models to (2-player) Stag Hunt-Games. Since a stag hunt has two Nash Equilibria (where one represents an efficient social contract while being risk-dominated by the other less efficient one) the question arises when evolution will lead to the more efficient though riskier outcome. This happens, they argue, if one allows for correlation entering during migration phase. Huttegger, concluded that, for stag hunts, introducing haystacks can help, and more generally that introducing structure in terms of group selection models to the populations, though it does not help altruism, often helps stag hunting, and always helps fairness.

All talks were accompanied by lively discussions and the conference took place in a relaxed and very communicative atmosphere. Noteworthy is also the rump session (where spontaneous speakers got five minutes sharp for giving a mini-presentation on what s/he is currently working on), which was fun.

Everyone interested in what else was going on at the Formal Ethics conference is highly recommend to visit iTunesU, where Video-Podcasts of all talks are available for download from the MCMP-Channel as usual. Extended abstracts and the whole schedule can be found on the website. The conference was generously supported by the Alexander v. Humboldt Foundation.

The new Formal Ethics Steering Committee, consisting of Albert J.J. Anglberger, Constanze Binder, Conrad Heilman and Paul McNamara, will coordinate future organisation and guarantee the continuity of the conference.

<div align="right">

Sebastian Hengst
Philosophy, LMU Munich

</div>

What If? 25–27 October

Is the ability to raise *'What if...?'* questions and suggest answers to them an essential feature of reasoning? This question itself can be stated counterfactually, as if to prove that counterfactual thinking really is all-pervasive: *What if* we had not learned to think about counterfactual scenarios? What if our brains had not allowed us to advance unconfirmed hypotheses and deduce consequences from those hypotheses? Would we not have been able to develop a rigorous system of knowledge (like modern science) or invent powerful works of fictions (like *The Odyssey*)? Would we be able to engage in any kind of abstract reasoning at all?

Over 40 conference participants from Europe and North America agreed that many precise answers to such 'iffy' questions feel rather speculative, hovering somewhere in the limbo between facts and fiction. The problems with counterfactual thought are legion. For example, when we raise these kinds of questions, it is unclear which assumptions about history, our nature and dispositions, etc. are held fixed and which are not.

But we can learn more about counterfactual thought when we look at its evolution. Peter Gärdenfors, the conference's keynote speaker from the Karolinska Institute in Lund, argued that while the development of the prefrontal cortex seems like a neural prerequisite for counterfactual reasoning, the behavioral correlate for the emergence of such reasoning can be found in activities such as playing games of make-believe, or

simply cooperating with peers and using tools to solve problems. Gärdenfors claimed that their highly developed problem-solving and tool-making behavior suggests that we have to ascribe some protoconcept of 'If...then...' to certain birds (e.g., crows) and higher mammals like chimps and dolphins.

Gärdenfors was also one of several participants of the conference who rejects the 'orthodox' philosophical approach to analyzing counterfactual claims, namely postulating a set of possible worlds and assigning truth values to counterfactual claims based on what happens in other, merely possible worlds that are said to be *further* away from or *closer* to our world according to how similar they are to our world.

Julian Reiss from the University of Durham echoed these concerns and argued that a more useful model for assessing counterfactual dependencies in the social sciences (like the one expressed by the counterfactual 'If Chamberlain hadn't been Prime Minister, the policy of Appeasement would have been abandoned much earlier') is inserting the event described by the antecedent into a causal structure, i.e., a framework that relates events to other actual and merely possible events in a probabilistic way.

Other scholars focused on more mundane aspects of counterfactual thinking. Historian Roland Wenzlhuemer investigated a deficient form of reasoning and theorizing that is very close to counterfactual thought but tainted with paranoia: Constructing conspiracy theories, often intended as alternatives to some 'official story', alternate accounts that seem to offer a real explanation of a complex, messy and confusing state of affairs that believers of conspiracy theories often have a hard time making sense of. Wenzlhuemer argued that conspiracy theories can be viewed as unintended caricatures of more respectable scientific theories.

The conference started with a reflection on the biological and cognitive evolution of counterfactual reasoning (Gärdenfors) and it ended with a reflection on the cultural and intellectual significance of such reasoning within certain currents of European modernism. Karin Krauthausen from Humboldt-University in Berlin focused on the concept of thought experiment, often attributed to Austrian physicist Ernst Mach in his 1897 Essay 'Über Gedankenexperimente'.

Krauthausen showed that Mach's ideas about thought experiments, while harshly criticized by his contemporary Pierre Duhem, are surprisingly close to the concept of psychic experimentation developed by a famous countryman of Duhem's, the author and poet Paul Valéry.

Thus, one major insight from the conference is that even though counterfactual thought and artistic imagination often seen like estranged siblings, both share a common origin. And time and again, after having traveled along their separate trajectories of intellectual history for a while, they seem to cross paths.

<div align="right">

Johannes V. Schmitt
Philosophy, University of Konstanz

</div>

Calls for Papers

Grammatical Inference: special issue of *Machine Learning*, deadline 1 December.

Weighted Logics for AI: special issue of *International Journal of Approximate Reasoning*, deadline 15 December.

Conformal Prediction and its Applications: special issue of *Annals of Mathematics and Artificial Intelligence*, deadline 15 December.

Hyperintensionality: special issue of *Synthese*, deadline 1 March.

What's Hot in . . .

Logic and Rational Interaction

The last months have seen some new articles on the Stanford Encyclopedia that might be of interest to the LORIweb community.

Pieter Adriaans presents a well written and concise account of the concept of information. He gives an overview of the conceptual history of information, starting from Plato up to the present. Adriaans presents the different meanings the term information can take, both in general use as well as in specific fields such as logic, physics and linguistics. In addition, he gives an overview of modern philosophy of information, making reference to Popper's falsificationist accounts as well as to Shannon and Kolmogorov.

In a second new entry Catarina Dutilh Novaes presents medieval theories of consequence. She argues that our conceptual intuitions behind modern consequence relations have already been formed in the middle ages and that consequently understanding the history of the consequence relation can give insights in why our intuitions are the way they are. Notably, there is a difference between syllogism and the consequence relation, though the exact scope of this difference has been debated. The entry gives an overview of theories of consequence and their development throughout the 13th and 14th centuries.

Besides that, a number of articles relevant to the community underwent major revisions, such as causal decision theory, formal representation of belief and combinatory logic.

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 Rasmus Rendsvig, our web manager or to the loriweb address.

<div align="right">

Dominik Klein
TiLPS, Tilburg University

</div>

Uncertain Reasoning

On 21 October 2012 six Italian seismologists were sentenced to six years of prison for multiple manslaughter. They were all sitting in the "Commissione grandi rischi" (Serious Risks Commission) when an earthquake killed more than three hundred people in the area of L'Aquila.

The sentence had massive international coverage and was received by many commentators as outrageous, if not outright medieval. The Guardian titled "A chilling verdict in L'Aquila". The Economist subtitled "In Italy, sloppy seismology can lead to prison". What is likely to be the shared view of many academics is articulated by David Spiegelhalter in his Understanding Uncertainty Blog. Nature devoted the 23 October 2012 editorial to "Shock and law".

On reading the above, and other pieces on this sensational verdict, I was often under the impression that many commentators failed to give due relevance to the judge's motivation. The scientists have been found guilty of misleading the public, not of having failed to predict the earthquake. They have been held accountable for their civil responsibilities and not for the accuracy of their scientific activity. One key passage in the judge's motivation refers to the fact that the scientists have supplied "imprecise, incomplete and contradictory information on the potential consequences ("pericolosità") of the seismic activity thereby hampering the procedures aimed at keeping the population safe".

The point is important, especially in the light of the enormous political pressure under which the "Commissione" was operating in 2009. Many political commentators in Italy suggest that the forthcoming appeal trial(s) will be heavily centered on the political motivation underlying the dramatically misleading overconfidence with which people in L'Aquila were told that it was safe to stay in their homes.

It isn't hard to predict that the appeal will be followed very closely by the international scientific community. One positive aspect of this, I think, is that it will provide a platform for discussion on many pressing issues including the quantification of severe uncertainty and the communication of the ensuing risks through the noisy channel which links experts, policy-makers and the general public.

HYKEL HOSNI
Scuola Normale Superiore, Pisa
CPNSS, LSE

INTRODUCING . . .

Managing Severe Uncertainty

Many of the important decisions that need to be made by individuals, groups or institutions are made under conditions of considerable uncertainty about the factors upon which these decisions depend. There are three essential, and interacting, dimensions to this uncertainty:

(i) Scientific or modelling uncertainty: we don't know what the state of the world is and under which dynamical law it will evolve.

(ii) Option uncertainty: we don't know what our options are and what would happen if we were to exercise one or another of them.

(iii) Ethical uncertainty: we don't know what value to attach to the possible consequences of our actions.

In the decision sciences uncertainty is typically handled by assigning probabilities to the various contingencies faced by policy makers and utilities to the possible outcomes of the decision. These probability and utility assignments are supposed to reflect, respectively, the information the decision maker holds about the state of the world and the consequences of any policy interventions, and what value she attaches to these outcomes. In conditions of severe uncertainty the decision maker cannot do this because for one reason or another she does not have the information, tools or resources to establish determinate probabilities for all relevant contingencies or to assign determinate utilities to all outcomes.

A salient example of severe uncertainty of this kind is that faced by climate change policy-makers. On the side of scientific uncertainty, they face the problem that climate models are imperfect and non-linear, which undermines standard forecasting procedures. In particular, climate models are known to be strongly imperfect in the sense that they do not even approximately mirror all the relevant features of real-world target systems, and this leads to the breakdown of probabilistic forecasts. On the side of option uncertainty, they face the problem that little is known about the potential effects (and side-effects) of some of the proposed measures for mitigation and adaption—carbon capture and storage, creating clouds from seawater, placing reflectors in the outer atmosphere, and so on—or what other mechanisms or technologies might be developed. Finally on the side of ethical uncertainty they face the problem of how to trade-off characteristics and effects of such interventions on different people at different times.

The problem of how to manage severe uncertainty will be the subject of an AHRC funded 3-year project, based in the Centre for the Philosophy of Natural and Social Science at the London School of Economics that starts in March 2013. The project team consisting of Richard Bradley, Roman Frigg, Katie Steele, Charlotte Werndl and Alex Voorhoeve will be supplemented by a Postdoctoral Research Fellow and two funded PhD students (for applications to these posts see here).

Taking climate science and policy as the main case study, the group will look mainly at two different families of questions.

1. *What are the reasons for scientific uncertainty regarding climate change?* This involves questions such as: Why and how do probabilistic forecasts break down? What model imperfections 'destroy' predictive power? How might we detect these imperfections? How can models be improved?

2. *How should one make policy decisions under uncertainty?* This involves questions such as: How can we make decisions when we lack determinate probabilistic predictions for crucial variables, such as local precipitation at future dates, upon which our decisions depend (this concerns scientific uncertainty)? How should we think about our possible policy options when we don't know which of them will be feasible (this concerns option uncertainty)?

Our aim is to develop answers to these two sets of questions that are sensitive to the way they interact. The manner in which we deal with scientific uncertainty, for instance, should be informed by the kind of information and level of accuracy required

in order to make policy decisions. Correspondingly, policy decisions should be made in a way which recognises their sensitivity to judgements that are themselves highly uncertain and in a way which maximises their robustness in the face of different possible resolutions of this uncertainty. Two workshops on these questions are envisaged and we would welcome ideas and contributions to them.

<div align="right">

RICHARD BRADLEY
Philosophy, LSE

</div>

EVENTS

DECEMBER

THE ANALYSIS OF THEORETICAL TERMS: Munich, Germany, 1 December.

LENLS 9: Logic and Engineering of Natural Language Semantics, Miyazaki, Japan, 1–3 December.

KNOWLEDGE ATTRIBUTIONS: Centre for Logic and Analytic Philosophy, KU Leuven, Belgium, 3 December.

NIPS: Neural Information Processing Systems Conference and Workshops, Nevada, USA, 3–8 December.

MM2012: Models and Mechanisms, TiLPS, Tilburg, Netherlands, 6–7 December.

K-NMTD: Konstanz-Naples Model Theory Days, University of Konstanz, Germany, 6–8 December.

CPH-LU: 5th Copenhagen Lund Workshop on Social Epistemology, Lund University, 7 December.

BAYESIAN OPTIMIZATION AND DECISION MAKING: Nevada, USA, 7 December.

MLINI: 2nd Workshop on Machine Learning and Interpretation in NeuroImaging, Nevada, USA, 7–8 December.

RISK AND ACCEPTABILITY: University of Zurich, 7–8 December.

PROBABILISTIC PROGRAMMING: Foundations and Application, Nevada, USA, 7–8 December.

25 YEARS IN CONTRADICTION: University of Glasgow, 7–9 December.

PROBABILISTIC NUMERICS: Nevada, USA, 8 December.

AGI12: 5th Artificial General Intelligence Conference, University of Oxford, 8–11 December.

AGI-IMPACTS: 1st Conference on Impacts and Risks of Artificial General Intelligence, University of Oxford, 10–11 December.

ICMLA: 11th International Conference on Machine Learning and Applications, Florida, USA, 12–15 December.

KPPC: The Philosophy of Science: How Seminal Works in Research Methodology Shape Our Understanding of Science and the Interpretation of Results, Prishtina, Kosovo, 14–15 December.

BELIEF CHANGE IN SOCIAL CONTEXT: University of Amsterdam, 14–15 December.

EGACRIS: Conference on Epistemic Groups and Collaborative Research in Science, Nancy, France, 17–19 December.

INTERNATIONAL TRIENNIAL CALCUTTA SYMPOSIUM ON PROBABILITY AND STATISTICS: Kolkata, West Bengal, India, 27–30 December.

JANUARY

SODA: ACM-SIAM Symposium on Discrete Algorithms, New Orleans, Louisiana USA, 6–8 January.

LFCS: Symposium on Logical Foundations of Computer Science, San Diego, California, USA, 6–8 January.

TARK: 14th Conference on Theoretical Aspects of Rationality and Knowledge, Chennai, India, 7–9 January.

ICLA: 5th Indian Conference on Logic and its Applications, Chennai, India, 10–12 January.

A&N: Aims and Norms: Reasoning, University of Southampton, 18 January.

CGCotPoM&L: 6th Annual Cambridge Graduate Conference on the Philosophy of Mathematics and Logic, Cambridge University, 19–20 January.

ULTRA-COMBINATORICS: Pisa, Italy, 24–25 January.

FEBRUARY

ICIIN: 2nd International Conference on Intelligent Information Networks, Maldives, 2–3 February.

SPIM: Workshop on Semantic Personalized Information Management, Rome, Italy, 4 February.

LAFLANG: 2nd International Workshop on Learning, Agents and Formal Languages, Barcelona, Spain, 15–18 February.

ICAART: 5th International Conference on Agents and Artificial Intelligence, Barcelona, Spain, 15–18 February.

CSEE: 2nd International Conference on Advances in Computer Science and Electronics Engineering, New Delhi, India, 23–24 February.

MARCH

METAPHYSICAL VIRTUES: Western Michigan University, Kalamazoo, Michigan, 15–17 March.

INFORMATION: 5th Workshop on Philosophy of Information, University of Hertfordshire, UK, 27–28 March.

UNILOG: 4th World Congress and School on Universal Logic, Rio de Janeiro, Brazil, 29 March–7 April.

APRIL

SBP: International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction, UCDC Center, Washington DC, USA, 2–5 April.

LATA: 7th International Conference on Language and Automata Theory and Applications, Bilbao, Spain, 2–5 April.

AISB: 6th AISB Symposium on Computing and Philosophy: The Scandal of Computation—What is Computation?, University of Exeter, 2–5 April.

The Analysis of Theoretical Terms: Munich, Germany, 3–5 April.

UNILOG: 4th World Congress on Universal Logic, Rio de Janeiro, Brazil, 3–7 April.

IMLA: 6th Workshop on Intuitionistic Modal Logic and Applications, Rio de Janeiro, 3–7 April.

ICANNGA: 11th International Conference on Adaptive and Natural Computing Algorithms, Switzerland, 4–6 April.

Information: Space, Time, and Identity: Milton Keynes, 8–10 April.

Models & Decisions: 6th Munich-Sydney-Tilburg Conference, Munich, 10–12 April.

Identity and Paradox: Lille, France, 11–12 April.

PAKDD: 17th Pacific-Asia Conference on Knowledge Discovery and Data Mining, Gold Coast, Australia, 14–17 April.

GCTP: Graduate Conference in Theoretical Philosophy, Groningen, Netherlands, 18–20 April.

Implicit Bias: University of Sheffield, 20–21 April.

SOoSI: The Social Organization of Scientific Inquiry, Center for Philosophy of Science, University of Pittsburgh, 20–21 April.

Explanatory Power: Understanding Through Modeling. Epistemology, Semantics, and Metaphysics of "Inadequate," Ruhr-Universität Bochum, 25–26 April.

NU/NDGC: 4th Annual Northwestern / Notre Dame Graduate Epistemology Conference, University of Notre Dame, South Bend, IN, 26–27 April.

AISTATS: 16th International Conference on Artificial Intelligence and Statistics, Scottsdale, AZ, USA, 29 April–1 May.

Courses and Programmes

Courses

BFAS: Spring School on Belief Functions Theory and Applications, Carthage, Tunisia, 20–24 May.

APhil: MA/PhD in Analytic Philosophy, 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.

MASTER PROGRAMME: Philosophy of Science, Technology and Society, Enschede, the Netherlands.

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 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 Sebastian.

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 MATHEMATICAL LOGIC AND THE THEORY OF COMPUTATION: Mathematics, University of Manchester.

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 Sebastian).

Open Mind: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.

PhD School: in Statistics, Padua University.

Jobs and Studentships

Jobs

Assistant Professor: in Logic or Analysis, Department of Mathematics, University of Connecticut, until filled.

Post-doc Position: in Artificial Intelligence, Institute for Artificial Intelligence, University of Georgia, until filled.

Post-doc Position: on Data Analysis for Knowledge Discovery and Decision Making, Department of Electrical, Computer, and Systems Engineering at Rensselaer Polytechnic Institute (RPI), Troy, NY, until filled.

Associate Professor or Professor: in Logic and the Philosophy of Science, University of Calgary, until filled.

Post-doc Position: in Probabilistic Reasoning, Vienna University of Technology, Austria, until filled.

Post-doc Position: in Cognitive Psychology and/or Computational Modelling at the Center of Experimental Psychology and Cognitive Science, Justus Liebig University Giessen, until filled.

Assistant Professor: in Cognitive Psychology, Center of Experimental Psychology and Cognitive Science, Justus Liebig University Giessen, until filled.

Post-doc Position: in Graphical Models / Structural Learning, Uncertainty Reasoning Laboratory, Queens College / City University of New York, until filled.

Post-doc Position: in Artificial Intelligence / Biomedical Informatics, Stevens Institute of Technology, until filled.

Post-doc Position: to contribute to the AHRC funded research project Managing Severe Uncertainty, The Centre for Philosophy of Natural and Social Science, LSE, deadline 1 December.

Assistant/Associate Professor: in Mathematical Logic, Institute for Logic, Language and Computation (ILLC), University of Amsterdam, deadline 3 December.

Lecturer: in Statistics, Northumbria University, deadline 6 December.

Post-doc Position: in Intelligent Robotics, Aalto University School of Electrical Engineering, deadline 7 December.

Research Assistants: in Machine Learning and User-Modelling for Human-Robot Interaction, Imperial College London, deadline 9 December.

Post-doc Positions: in Philosophy of Mind/Language working on "The Nature of Representation" project, Philosophy, University of Leeds, deadline 19 December.

Professor: in Statistics, University of New South Wales, Australia, deadline 13 January.

Post-doc Position: in Metaphysics of Science, Institut d'Histoire et de Philosophie des Sciences et des Techniques, Paris, deadline 15 February.

Studentships

PhD Position: on the project "Models of Paradox," Philosophy, University of Otago, until filled.

PhD Position: on Data Analysis for Knowledge Discovery and Decision Making, Department of Electrical, Computer, and Systems Engineering at Rensselaer Polytechnic Institute (RPI), Troy, NY, until filled.

PhD Positions: in the Statistics & Probability group, Durham University, until filled.

PhD Position: in Intelligent Robotics, Aalto University School of Electrical Engineering, deadline 7 December.

PhD Position: in Logic-Based Decision Support, Individual and Collective Reasoning Group, University of Luxembourg, deadline 15 December.

PhD Positions: in Theoretical and Computational Neuroscience and Machine Learning, Gatsby Computational Neuroscience Unit, University College London, deadline 16 December.

PhD Position: on research project "Epistemic Utility Theory: Foundations and Applications," Philosophy, University of Bristol, deadline 20 December.
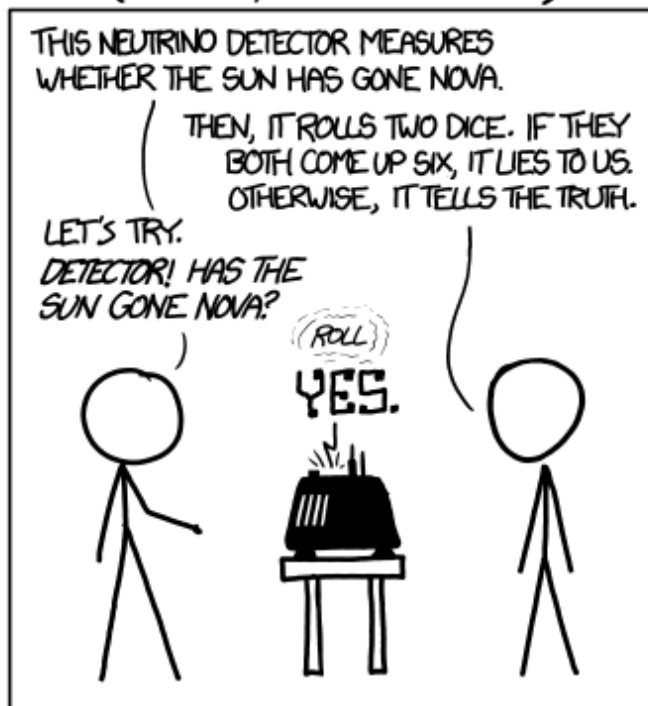
PhD Position: on the project "Knowledge Representation and Inference Based on Type-2 Fuzzy Sets and Systems," School of Computer Science, University of Nottingham, deadline 30 December.

Two PhD Positions: for research project on "Managing Severe Uncertainty," Department of Philosophy, Logic and Scientific Method at the London School of Economics and Political Science, deadline 11 January.
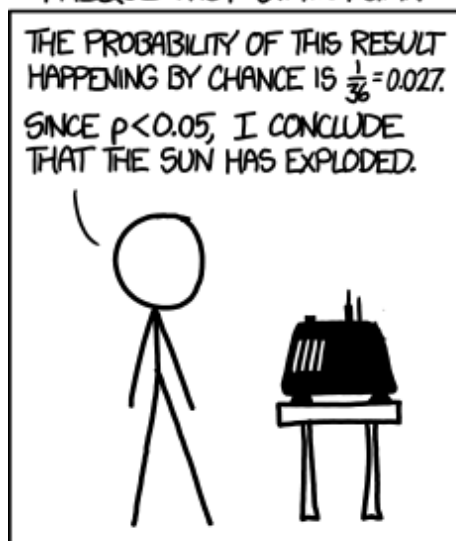
PhD Position: at the Institute for Logic, Language and Computation (ILLC), University of Amsterdam, deadline 14 January.

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