THE REASONER

CONTENTS

EDITORIAL

Over the last decades, the relations between psychology, logic and philosophy on the topic of reasoning have been somewhat strained. Once the Piagetian idea that classical, traditional logic provided a good descriptive model of human reasoning was thought to have foundered on Wason's critique in the 1960s, psychologists were increasingly convinced that logic had nothing to offer for their investigations. Logicians, in turn, were developing increasingly sophisticated and varied logical systems, but paying little attention to how (and whether) these systems were in fact connected to human reasoning *au naturel*. Philosophers, at the other end, held on to the idea that philosophy is about what *ought* to be the case, not about what *is* the case, and thus for the most part viewed the results in the psychology of reasoning as philosophically irrelevant.

The absence of contact between these different disciplines is still essentially the current situation, but with a few notable and very exciting exceptions. The work of Keith Stenning in particular boldly breaks down the borders between these disciplines, and thereby offers a much more encompassing and nuanced account of human reasoning than what can be found elsewhere.

Keith is an honorary professor at the University of Edinburgh, where for years (1990–2000) he led the Hu-

man Communication Research Center. He is the author of countless articles and two books: *Seeing Reason* (OUP, 2002) and, with Michiel van Lambalgen, *Human Reasoning and Cognitive Science* (MIT Press, 2008). He is now retired, but fortunately for us all, has remained an extremely active researcher—in fact, I am glad to report that Keith and I are currently developing a research project to investigate the effects of schooling on reasoning by means of cross-cultural experiments in Europe and Africa.

What follows is a condensed version (for reasons of space) of the conversation I had with Keith in Groningen in March of 2012; the actual interview ended up being much longer and covering a wide range of topics, because Keith has fascinating ideas on pretty much any topic one can think of! He talks about the status of cognitive science and psychology as disciplines, his interdisciplinary professional trajectory, and his recent collaboration with Michiel van Lambalgen tracking down the inherent component of defeasibility in human reasoning.

<div align="right">

CATARINA DUTILH NOVAES
Philosophy, Groningen

</div>

FEATURES

Interview with Keith Stenning

Catarina Dutilh Novaes: Keith, thank you for agreeing to do this interview! You have very broad interests, ranging from reasoning and cognition to language and evolution, amongst others. So how do you like to be identified professionally? Are you a cognitive scientist, or a psychologist, or a philosopher?

Keith Stenning: I generally think of myself as a cognitive scientist, and I like that label because I regard it as a label above the disciplinary level. So I don't think of cognitive science as another discipline, in fact I like to think of it actually as a rejection of disciplinary boundaries. I wouldn't call myself a philosopher because I am interested or involved with really a very small part of philosophy. I was always a philosopher who was mostly interested in how philosophy can influence science, psychological science, social science.

CDN: Cognitive science being a broad label, that would only make sense if you understand cognition there in a very broad sense, which I know you do. But perhaps you would like to elaborate on this a little bit.

KS: Well again, there are those who take cognition to mean the computation of exact problems or something like that, and they think of it as what is sometimes described as 'cold cognition'. And I see that as a mistake: those may be the places you start when you're trying to make progress with a horrendously large topic, but that's definitely not the goal.

CDN: By the way, if you could enlighten us on the difference between a cognitive scientist and a psychologist that would be very helpful, since I still don't know myself!

KS: Cognitive psychologists are those who for the most part restrict themselves to just doing experiments, which makes it a very bizarre science. Sciences don't consist of just doing experiments. But the way the disciplines are carved up in the social sciences, they are mainly carved up methodologically, so I do experiments and I'm a psychologist, and you do theory, so you must be doing something else. Maybe it's AI, maybe it's computation, maybe it's philosophy, maybe it's mathematics.

But there are cognitive psychologists who are deeply interested in what goes on in philosophy, what goes on in AI, computation, such as John Anderson. He uses AI techniques to simulate data that he gets from experiments. I think that's a paradigmatic example.

CDN: So the idea would be that the psychologist is really almost exclusively focusing on the 'doing experiments' level of investigation, whereas a cognitive scientist would also include the experimental data but wants more.

KS: Psychologists think they have theories. On closer examination it's not what any other science means by a theory. I am methodologically basically a psychologist. Most of what I do and publish in the end is psychological data. It's extraordinarily hard, it's at a very early stage, and biology went on like this for centuries. But biology had some very helpful informal theories, like taxonomy, and anatomy, and again these are proto-theories of one sort or another. So I have a lot of sympathy with the difficulty, but I have no sympathy whatsoever with somebody who thinks they don't need

theory.

CDN: So you don't resent them for not having theories yet, but rather for thinking that they don't need theories to start with.

KS: And what you will inevitably find is that they have invented them implicitly in the back door, very often they are logical theories, and they are used very badly, because they are never acknowledged.

CDN: So what about your philosophy background, and your professional trajectory?

KS: I was an undergraduate, doing a joint double honours degree in philosophy and psychology, at Oxford. Two groups of people who thought their two subjects had nothing to do with each other had been forced to teach this way, because psychology was just starting and the founding fathers deemed that it was a dangerous new subject and they needed to keep their feet on the ground. The only town in the world where philosophy, it was thought, might keep your feet on the ground. . .

But actually before I was a psychologist or a philosopher I was a biologist. I had a wonderful biology teacher in the last two years of secondary school. And I still think of myself sometimes, and in fact increasingly as I get older, as having an important biological input.

CDN: Yes that's definitely clear from your work, biology is always in the background. So what happened next? You were a philosophy and psychology undergraduate, then what?

KS: Well, being an awkward sod, I was dedicated to the idea that these two subjects were rather close. And here were these two groups of people who thought they had nothing whatsoever to do with each other, so I spent my time winding up the lecturers, and I learnt a lot, I learnt an awful lot. It's not a bad strategy for a student, although not one they can pursue very easily nowadays.

CDN: And where did you do your PhD?

KS: I went west young man: I loved Oxford intellectually but found it socially an impossible place, so I went off to the New World. I went to do a PhD with George Miller in Rockefeller University, because he had then been doing early work with Chomsky on what it is to know a language. It was a wonderful environment, where they had bought themselves a philosophy department, rather in the spirit with which one buys a basketball team! So I could do both psychology and philosophy.

I wrote a terrible, long, straggly mess of a thesis and George was kind enough to get me a PhD on the basis of it. At the time there were wonderful defeasible logics just sprouting, they were the first ones, but there were also the first studies showing they were wildly intractable.

CDN: Which is a paradox right? Defeasible logics are supposed to be the logics underlying our reasoning; we can reason, and yet, these logics are intractable.

KS: The gleam in the eye was that they would make things easier, and it was rapidly shown that they make things 100 times worse than classical logic. So that was all going on in the background, and with my very minimal technical understanding anyway, I was confused about whether they were logics.

CDN: So after your PhD you went back to the UK?

KS: I did a postdoc in George's lab and then I went back to Liverpool, to the psychology department in Liverpool. And then I went to Edinburgh. I've been in Edinburgh ever since. From 1990 to 2000 I was heading the Human Communications Research Centre, a wonderfully interdisciplinary place; the ideal situation for me.

CDN: So now going forward quite a bit, what you have been doing in recent years is mostly collaborating extensively with my former colleague Michiel van Lambalgen, and your joint 2008 book (*Human Reasoning and Cognitive Science*) is in many respects groundbreaking. Personally, I think you argue convincingly that the way psychologists of reasoning have been thinking about reasoning is all misguided, and the same holds of philosophers in their thinking about logic and reasoning. Perhaps you would like to tell us a bit about how this collaboration came about and how these groundbreaking ideas started to emerge.

KS: It was Michiel who had the insight to recognize that we needed to work together. I was in Amsterdam for something or other. I think I probably gave a talk and he came up to me and he said he was a mathematician and a probabilist, but he was really interested in cognition, and modeling cognitive processes like reasoning. He'd seen some of the stuff that I'd done, and he thought that defeasible logic was needed. And the defeasible logic that he knew had changed, I hadn't really been tracking what was going on.

CDN: So the monstrously complicated systems of a few decades before had been improved, in the meantime?

KS: Right. So Michiel had all the technical skills that I lacked and I had the cognitive questions coming out of my head at a rate faster than even he could solve them. I mean here was somebody who could offer a phenomenon and maybe some data, and he would come up with a model and one would say "well, that's very

good here but doesn't seem to fit over here", and he would say, "oh well let's have a different model". So an amateur who struggled to grasp one particular kind of model, would stick with this model come what may; but a master of the art will simply say, "bugger that, let's start over", or modify it, or whatever.

We started on Wason's task. It's an incredibly small, encapsulated task, an innocent looking thing. And I'd known about this since I was an undergraduate student, I think, and it just seemed to me to be stupid. Everybody knew, I thought, that the natural language conditional, whatever it was, *wasn't* the material implication, and you needed the material implication interpretation in order to fall into Wason's trap. So I always refused to have anything to do with it. Apart from teach it: I used to teach it as a way not to do psychology. But when Michiel turned up, and had a logic which on the face of it makes it quite a good account of one of the many readings that people have for the natural language conditional, and a plausible first one to jump at in this task, then the ground kind of shifted.

CDN: And this plausible reading is the defeasible conditional right?

KS: It's a law-like one. I mean one knew from the philosophy of science what these things were like, although the philosophy of science at that point didn't have—to my knowledge—formalized accounts of what law-like conditionals were—but they were resistant to counterexamples, that's the most important thing.

CDN: So law-like conditionals, unlike the material conditional, allow for counterexamples. They're robust, in that sense. You and Michiel offered a formalized account of the phenomenon taking inspiration from the system which is known as *closed world reasoning*. Perhaps you could say something about that.

KS: Closed-world reasoning is the tractable form of defeasible reasoning that had arrived in the middle 90s. So we start a story: once upon a time there was a cat, so we've got a domain, with a cat in it, which interacts with the text as it comes in. But here you've got a logic that's so tractable that you've got a unique minimal model at every stage. And the propositional form, which is the form we use for modeling most of the psychological experiments, is extremely simple, so the unique model is simply a set of valuations of the incoming premises in the text, and everything looks like material implication until you hit a possible counterexample. So 'If Mary has an essay she'll be in the library'. But then maybe the library is closed. So if she won't be in the library now what's going to happen? So the closed world is this lit-

tle bit of the universe we're constructing the model of, and the crucial point of the closure is that we can reason about what is known about this. If we don't have any evidence of an abnormality on a particular matter, then we will assume there is no abnormality. That's what makes the whole thing tractable, and what makes the model unique and it gives it all the psychological properties. We can implement this in a neural network where the reasoning happens time-linear with the depth of a search in the network, which is extraordinarily efficient.

CDN: So there's a closed world assumption which is that, in the absence of evidence to the contrary, we assume that nothing abnormal is going on. This assumption is not only technically tractable, or gives rise to technically tractable systems, but you claim that it's also psychologically plausible, that that's exactly how people tend to reason.

KS: Yes. This database of knowledge that you have, this huge database, is organized so that you jump to conclusions all the time, and those conclusions may then be overridden by what happens next.

CDN: Which thus nicely captures the idea of defeasible reasoning: conclusions may be overridden by what happens next. I wish philosophers would understand that defeasible reasoning is the real thing! (laughs)

More Right Squares

In a previous paper (The Right Squares, *The Reasoner* 6.6:100–101) I showed that deontic and temporal logics are about distinct types of proposition, and only the traditional synonymy definition of propositions can support the required distinctions. This led to correcting the standard squares of opposition associated with them. Two further cases are where there are alethic or epistemic operators rather than deontic, or temporal ones.

For is it invariably the case that Mp ∨ M¬p, i.e., that it is possible that p or possible that not-p? If that was so then there would be an automatic equivalence between "It is not necessary that p" and "It is possible that not-p", i.e., between "¬Lp" and "M¬p", allowing no distinct use for "It is not necessary that p" with regard to propositions not in the right category. The alethic operators, however, are in truth no different from the deontic and temporal ones in this respect. For some propositions are ruled out by the assertion of "Mp ∨ M¬p"; they are the nonsensical ones which Chomsky called selection mistakes (see Hodges, W. 1977: *Logic*, Penguin, Harmonsworth. pp. 21–22). In Rylean terms, they are

the ones that involve category confusions like "My cat is a prime number" and "The moon is in E Flat Major". Such sentences can be negated, and also translated, so they express propositions on the "synonymy" view, but the propositions expressed are nonsensical. One can certainly say that the ocean was at one o'clock, but it does not make sense to say either that or its negation, since if it is nonsense that p ($\neg$Sp), then it is also nonsense that not-p ($\neg$S$\neg$p). Thus "Sp" is the equivalent of "Mp $\lor$ M$\neg$p", and "$\neg$Sp" is the equivalent of "$\neg$Mp & $\neg$M$\neg$p", and each may be asserted truly.

Standard Predicate Logic, in its abstraction, lists supposed predicates but does not consider which combinations of predicates make sense or nonsense, such as those deriving from the same or different category of predicate. It is not surprising in this context, therefore, that many operator logics have ignored categorical differences between propositions, such as those appropriate for deontic and temporal operators; and the associated operator in the alethic case, "It makes sense that p", as a result, has been almost completely overlooked.

The most significant consequence is that Modal Logic has continually ignored the proper, grammatical basis for possibility, thinking that it is never the case that $\neg$Mp & $\neg$M$\neg$p, and that "Lp" is equivalent to "$\neg$M$\neg$p". But "Lp" is equivalent to "$\neg$M$\neg$p & Sp", in a perfect square exactly like those defined in the previous paper:

A: $\neg$M$\neg$p & Sp,

E: $\neg$Mp,

I: Mp,

O: M$\neg$p $\lor$ $\neg$Sp.

So "Sp" identifies that select set of propositions whose respective truth and falsity determine the various possible worlds. "Sp" is, of course, then contingent on "p", but it is not empirical, being based on conceptual notions of compatibility. Writing "In world i it would be true that p" as "Tip", then *for these propositions* Mp $\equiv$ ($\exists$i)Tip, where (i)($\neg$Tip $\equiv$ Ti$\neg$p), and Tip $\equiv$ p when i is the actual world. Of course "Lp" then entails "p", since Lp $\equiv$ (i)Tip, and so in the actual world, p.

So what might be obligatory is some behaviour, what might always happen is some event, and what might be possible is something that makes sense. But the basic factor that links all these cases is the revealed facility in language to discriminate an external from an internal

negation. As Aristotle realized, the O form "Not all S is P" is distinct from the I form with a negated predicate "Some S is not P", unless there are definitely some S (see Thompson, M. 1953: On Aristotle's Square of Opposition, *Philosophical Review* 62, pp. 251–265).

It is this feature that has one final parallel, in the definition of Knowledge. It is sometimes said that Knowledge cannot be defined, and so finding a definition is perhaps rather surprising. But consider the following perfect square:

A: $\neg$($\exists$e)(pr($\neg$p/e) > 0) & ($\exists$e)(pr(e) = 1),

E: (e)$\neg$(pr(p/e) > 0),

I: ($\exists$e)(pr(p/e) > 0),

O: ($\exists$e)(pr($\neg$p/e) > 0) $\lor$ $\neg$($\exists$e)(pr(e) = 1).

Here e is the evidence on which a person is basing his knowledge of p, and for it to be indeed knowledge then some of that evidence must itself be true otherwise there are problems of the sort illustrated by E.L. Gettier (1963: Is Justified True Belief Knowledge? *Analysis* 23, pp. 121–123). Of course if there is some evidence that is true, the given square collapses in a parallel way to the others before. For then there is no distinction between the O form and an I form with an internal negation in it; that is to say, then either ($\exists$e)(pr(p/e) > 0) or ($\exists$e)(pr($\neg$p/e) > 0). But from the two conjuncts in the A form (excluding cases where there are infinite sets, since the concept of knowledge predates measure theory), there follows the truth of p itself, since there is no support for not-p, which means that the true evidence must totally support p. Also the I form can be read, "there is some reason to think that p", i.e., "it is believable that p". Hence, if the A form defines Knowledge, it becomes provable that Knowledge not only entails Truth, but also Belief. The role of Justification is then sufficiently explained in the relation between the evidence and the proposition known, making the traditional definition of knowledge as justified true belief almost correct. But the truth that must first be attended to is not the truth of what is known, and instead it is the truth of the evidence on which what is known is based.

On this account beliefs are not closed under logical consequence, since Adjunction does not hold for probability measures: p can be probable, and q can be probable, without the conjunction of p and q being probable. But a form of "belief revision" is ready to hand in terms of conditionalisation, since the link between beliefs and

probabilities is central to Bayesianism. It might seem, however, that the problem of Omniscience recurs with the above definition of Knowledge: don't all necessary truths have probability 1 based on any evidence whatsoever? Maybe so: but the question is whether Goldbach's Conjecture, for instance, is a necessary truth. Only after we had established it was a necessary truth would its probability be 1, and before that its probability is not settled.

HARTLEY SLATER
Philosophy, University of Western Australia

A Hypodox! A Hypodox! A Disingenuous Hypodox!

Frederica was born on the 29th February. While her twin, Frederic, is overseas negotiating with pirates about his current age, Frederica is wondering whether to celebrate her 21st birthday on the 28th February or 1st March. (Her preference is both, but her orthodox parents will not allow that. They say a birthday should be just one day a year, so she cannot have two birthdays.) Her American father suggests 28th February, as it is the last day of the month. Nevertheless, her English mother thinks it odd to celebrate before the 29th, so she believes 1st March is the best choice. Frederica can celebrate on either day consistently, but lacks a decisive reason. Her choice is underdetermined.

Gilbert's paradox is about Frederic being indentured to a pirate until he turns 21. Frederic has lived 21 years, but has only had 5 birthdays. It ranks about .1 on Sainsbury's Richter scale of paradoxes (where the Barber rates 1 and the Liar and Russell's are at 10). It has the above (equally weak) *hypodox* about when Frederica should celebrate her 21st birthday.

The Truth-teller is a paradigm hypodox. The Truth-teller may consistently be either true or false, but it seems there is no basis for determining which it is. Grelling's hypodox concerns whether 'autological' is autological. And Russell's hypodox concerns whether the set of sets that are members of themselves is itself self-membered.

There are even hypodoxes of time travel (Eldridge-Smith 2007 'Paradoxes and Hypodoxes of Time Travel', in *Art and Time*, Lloyd Jones, et al. eds., Melbourne: Australian Scholarly Publishing, 172–189). In the Grandfather paradox, Tim time travels to assassinate his grandfather, before the conception of the intervening parent. It seems mutually exclusive events would obtain in this scenario (without branching time). In the Grandmother hypodox, Tess has always wondered who saved her grandmother. She time travels to witness the event. Realising that no-one will snatch her 5 year old Gran from in front of the tram in time, she does it herself. No mutually exclusive events occur; nevertheless, it seems indeterminate whether this scenario is possible. Speculative fiction contains ample examples of these inconsistent and consistent scenarios. Lewis' *laissez faire* time travel may avoid paradoxes, but it allows such hypodoxes (Lewis, D. 1976 'The Paradoxes of Time Travel', *American Philosophical Quarterly*, 13: 145–52).

In contrast, the Sorites does not seem to have a hypodox. Moreover, given the relation between Russell's and Cantor's paradoxes, a hypodox for Cantor's is notable by its absence. Indeed, many paradoxes (e.g., Zeno's, Newcomb's and Hempel's paradoxes) do not seem to have hypodoxes at all.

Hypodoxes are characterized by the lack of some principle to determine which among an inconsistent set of individually possible or plausible options is correct. (They are not due to ignorance of contingent information.) Hypodoxes include whether the Truth-teller is true or false, and whether the Grandmother scenario is possible. The set of all self-membered sets cannot consistently both be and not be a member of itself; it could consistently be either, yet there is a lack of a principled determination.

Intuitions vary. Some killjoys, fundamentalists and pirates maintain that Frederica cannot celebrate a birthday on either 28th February or 1st March, but must wait for the next leap day! There is an analogous intuition that the Truth-teller is neither true nor false. Mortensen & Priest (1981 'The Truth Teller Paradox', *Logique et Analyse*, 24: 381–388) derive a paradox from this alternate intuition by an ingenious, surprising argument. The principle of sufficient reason (or the Truthmaker principle) might seem to support this alternate intuition were it not for the resulting contradiction.

Nevertheless, semantic and set-theoretic hypodoxes do not need the schemata for conceptual principles that characterize their paradoxes. The relevant instances of such schemata are derivable as logical theorems. For example, Russell's hypodox has no need of Abstraction, because the relevant instance of Abstraction is a theorem:

(4) The set of self-membered sets is a member of itself iff the set of self-membered sets is a member of it-

self (by the Identity Theorem of Sentential Logic)

Furthermore, the relevant instance of Comprehension is a theorem:

(5) (∃*y*) (The set of self-membered sets is a member of *y* iff the set of self-membered sets is a member of itself) (from (4) by ∃I)

It seems paradoxes use principles that their hypodoxes do not require. Contrariwise, hypodoxes seem to be proto-paradoxes; by adding principles, a hypodox may become a paradox. Consider the pair of sentences:

(6) The other sentence in this pair is false.

(7) The other sentence in this pair is false.

Consistent answers are that (6) is true and (7) false, or (6) is false and (7) is true. They may seem at first hypodoxical, as nothing determines which is true and which false, but they could consistently be either. However, if we add a "symmetry" principle to the effect that tokens of the same sentence type, such as (6) and (7), have the same truth value in the absence of contextual differences, then we have a paradox (Sorensen, R. 2001. *Vagueness and Contradiction*, Oxford: Clarendon Press, p. 166).

The concept of a hypodox as a consistent kind of conundrum akin to a paradox helps frame a number of research questions. E.g.,: Do all hypodoxes have associated paradoxes? What is the relation between a hypodox and its paradox? Which paradoxes have associated hypodoxes?

One more example, the Eldridge-Smith paradox (the ESP), is a truth-functional relative of the Liar (Eldridge-Smith 2008 *The Liar Paradox and its Relatives*, Australian National University):

(8) My favourite biconditional is 'My favourite biconditional is not true iff *Q*'

If *Q*, then the biconditional referred to in (8) is paradoxical, and if ¬*Q*, then it is hypodoxical. Furthermore, compare (8) with (9):

(9) My second favourite biconditional is 'My second favourite biconditional is true iff *Q*'.

When (8)'s biconditional is hypodoxical, (9)'s is paradoxical; and when (8)'s is paradoxical, (9)'s is hypodoxical.

Peter Eldridge-Smith
Philosophy, ANU

News

Agent-Based Modeling in Philosophy, 15–19 May

On May 15–19 the first Ghent-Tilburg workshop on agent-based modeling in philosophy took place in Spa, Belgium. Simple agent-based models can produce surprisingly strong results. 43 years ago Thomas Schelling demonstrated with his famous checkerboard model that a small racial preference is sufficient to produce strict segregation over time. In recent years agent-based modeling has become common currency in philosophy with work on the emergence and change of moral behavior (Alexander (2007) *The structural evolution of morality*, Cambridge University Press), the social structure of science (Zollman (2007) The Communication Structure of Epistemic Communities, *Philosophy of Science* 74(5), 574–87) and the division of cognitive labor (Weisberg & Muldoon (2009) Epistemic Landscapes and the Division of Cognitive Labor, *Philosophy of Science* 76(2), 225–52). Easily accessible but powerful software (NetLogo) has been developed to democratize agent-based modeling. Yet many philosophers lack even the most basic programming experience and do not know where to begin. The Ghent Center for Logic and Philosophy of Science and its doctoral schools and the Tilburg Center for Logic and Philosophy of Science (TiLPS) joined forces to give participants a basis to start building their own agent-based models. In order to build, apply, or interpret a computational model, certain basic skills are needed. These can be roughly divided into three categories that formed the basic tenets of the workshop programme: conception, coding and valorization.

In this five-day workshop morning sessions—led by Aaron Bramson—focused specifically on learning how to *code* Netlogo models. Afternoon sessions focused on conception and valorization for philosophical purposes. These sessions were led by Kevin Zollman and Ryan Muldoon, philosophers with particular expertise in the application of agent-based models to topics in philosophy. Agent-based models are simulation models that aim not so much at making correct predictions but rather at understanding what the essential features are and how they interact. This focus on the process rather than the solution prompts novel philosophical questions, for which agent-based simulation provides new answers. How does network structure affect the spread of knowledge in communities? What combination of strategies allows communities to maximize sci-

entific progress? How should communities trade off exploration against exploitation? A key to the *conception* of an agent-based model is to formulate a clear research question. Every model needs a purpose, and this purpose will inform modeling decisions later on. Once the model is up-and-running, publishing the results (*valorization*) is the next challenge. Because agent-based modeling is such a new approach in philosophy, few conventions have developed about how to present an agent-based model in a journal article. It was therefore very useful to have experienced philosophical modelers give an overview of what they consider good practice for reporting modeling results in philosophical journals.

<div align="right">

ROGIER DE LANGHE
TiLPS, Tilburg University

</div>

Reasons and Rationality, 20–22 May

The third St. Louis Annual Conference on Reasons and Rationality (SLACRR) took place May 20–22 at the Moonrise Hotel in St. Louis, MO. The conference, sponsored by the Philosophy Departments at UM-St. Louis and Washington University, featured a keynote address from Jonathan Dancy (Texas) entitled, "More Right than Wrong." The organizers are already planning for the fourth installment of SLACRR. Abstracts for SLACRR 2013, which will feature a keynote address from Michael Smith (Princeton), will be due on January 1, 2013. The conference aims to be a forum for new work on practical and theoretical reason, broadly construed. More information is available here. Here are brief summaries of the some of the papers presented at SLACRR 2012:

In "Acting and Believing for Reasons," Nomy Arpaly (Brown) and Timothy Schroeder (Ohio State) defend the idea that thoughts and actions caused by reasons—in virtue of the fact that the reasons rationalize the thoughts or actions—are thoughts reached *for* a reason, or actions done *for* a reason.

In "Introducing Socratic Anti-Intellectualism," Agnes Callard (Chicago) argues that we can learn something about contradictions from the practice of Socratic refutation, namely, that they are more assertible than they are believable.

Patricio Fernandez (Harvard) presented a paper titled "Why not Act?" in which he defends the thesis that the conclusion of practical reasoning is an action, proposing to understand it in terms of certain acts of a rational subject, and addressed some common objections to it.

In his paper "Explaining Constitutive Norms," Ernesto Garcia (Massachusetts) discusses various requirements for being a rational agent—what they are and how, if at all, they might help to guide our actions.

In his paper "A Very Good Reason to Reject the Buck-passing Account," Alex Gregory (Reading) argues that the buck passing account of value is mistaken because reasons themselves can be evaluated as better or worse.

In his paper, "A Puzzle for Analyses of Rationality," Ali Hasan (Iowa) discusses a puzzle facing any account of rationality that allows for rational but false beliefs about one's own rationality, and argues that the best way out of the puzzle is to reject two apparently intuitive principles of rationality: (a) the principle that if it is rational for S to believe that the belief that p is rational for S then this belief is rational for S, and (b) the principle that if it is rational for S to believe that the belief that p is not rational for S, then it is not rational for S.

In his paper "Perform Your Best Option," Doug Portmore (Arizona State) argues whether an agent ought to perform some temporally discrete option (such as the option of typing the letter 't' just now) depends on whether the agent ought to perform some temporally extended option (such as the option of typing this sentence over the last minute) that involves performing that temporally discrete option.

In "Team Reasoning, Shared Agency, and Non-evidential Warrant for Belief" Abe Roth (Ohio State) argues that if, as its advocates maintain, team reasoning and rationality is distinct from individual instrumental rationality, then there must be a peculiar non-evidential warrant for when it is appropriate to employ this reasoning. Roth ties such a warrant with the ability to exercise the capacity for joint or shared agency.

Mike Titelbaum's (Wisconsin) paper "In Defense of Right Reasons" argues that any agent who has false beliefs about the requirements of rationality makes a rational error. This thesis follows surprisingly easily from the claim that akrasia is irrational, and has the consequence that agents who evaluate their evidence correctly should stick to their evaluation even in the face of a disagreeing peer.

In his paper, "Reasons for Belief, the Aim of Belief, and the Aim of Action," Daniel Whiting (Southampton) argues that the aim of belief is derived from the aim of action, and that this explains why subjects do not take

there to be practical or pragmatic reasons for believing.

John Brunero
Eric Wiland
Philosophy, UM-St. Louis
Charlie Kurth
Philosophy, Washington University

Experts and Consensus in Economics and the Social Sciences, 25–26 May

Carlo Martini, Marcel Boumans and Niels Gottschalk-Mazouz, with the support of a 'Deutsche Forschungs Gemeinschaft' (DFG) grant, organized a two-day workshop on 'Experts and Consensus in Economics and the Social Sciences' at the Center for Philosophy & Economics at Bayreuth University (Germany). From the 25th till the 26th of May several speakers presented their recent findings in the field and contemplated on future routes of research. A total of 12 talks were given, and the second day ended with a roundtable conversation, which allowed the participants to elaborate on future opportunities for continuing research and collaboration on the topics of the workshop.

On Friday morning Marcel Boumans kicked off the workshop by tackling the problem of rational consensus in economics: the focus was on the "Cooke method", a mathematical aggregation method of weighing the opinion of each expert on the basis of his or her knowledge and ability to judge relevant uncertainties. The method was compared with other consensus models (e.g., the Delphi method).

Filip Buekens argued why the notion of accuracy is important to keep in mind when talking about truth and expertise. In particular, he attacked Goldman's reliabilist view of knowledge, as merely focusing on reliability and truth, by urging for a third dimension, i.e., accurate beliefs, to be included.

Maria Jimenez Buedo addressed the issue of how to attribute expertise in uncertain times. She highlighted a dilemma with some objectivist notions of expertise: on the one hand, underplaying the relational aspect of expertise leads to a concept that is indistinguishable from knowledge; on the other hand, underplaying the objective sense of expertise leads to the stretching of the expert status.

Merel Lefevere presented joint work with Eric Schliesser, where they defend the thesis that the character of scientific communities can be evaluated morally and found wanting in terms of moral responsibility. By way of critical discussion of a recent proposal by Heather Douglas, they argued that even an epistemically successful scientific community can be held morally responsible for consequences (also unforeseen ones) that follow from policy advice given by its members.

Rafal Wierzchoslawski described the role of experts in the condominium model of republican (re-)solution of social, economic and political problems, drawing on insights from Turner's book 'Liberal Democracy 3.0'.

Frank den Butter ended the first day of talks by describing the institutional economics of stakeholder consultation, and arguing for a reduction in implementation costs through 'matching zones', which focus on bringing all relevant stakeholders together in an institutional setting. His arguments allowed him to emphasize the often overlooked difference between compromise and consensus.

Robert Evans opened the second day of the workshop, presenting the analogy of 'emperors, mavericks and children' as a way of analyzing problems of expertise in science and society. He stressed the necessity (and difference) of two phases in the decision making process: a technical phase (related to questions of fact) and a political phase (related to questions of preference).

Amir Konigsberg talked about disagreement: he provided an account of the difference between first-order (subjective) and second-order (objective) evidence, and explained how the current literature on disagreement fails when dealing with what he thinks is the real problem of disagreement.

Aviezer Tucker showed us how applying the Neyman-Rubin model of causal inference to the explanation of expert consensus can point out certain obscurities.

Carlo Martini presented an analysis of some normative principles of expertise with the case study of the Monetary Policy Committee. He suggested that formulating a number of principles to be followed when employing experts in a committee can go a long way towards optimal institutional design in committees dealing with economic issues.

Laszlo Kosolosky, in work together with Jeroen van Bouwel, tried to explicate ways of consensus making in science and society, and more importantly at its interface, by introducing a procedural account of consensus as to deal with the tension between consensus and plurality, resulting in a social account of consensus formation.

J.D. Trout excelled at the difficult task of closing a round of very interesting talks. His talk highlighted the difficulties intrinsic to the pretense of democratizing science and employing expertise in society, and opened some avenues to possible alternatives.

To conclude, we would like to thank all the participants in the workshop for the many interesting debates that took place during every Q&A session, as well as for the extremely friendly and jovial atmosphere during the workshop.

Laszlo Kosolosky
Philosophy, Ghent
Carlo Martini
Philosophy and Economics, Bayreuth
TiLPS, Tilburg University

A Priori Justification, 16–17 June

The conference was organised as the second and last major international conference of the AHRC-funded project on Basic Knowledge (2007–2012), led by C. Wright and hosted by the Northern Institute of Philosophy at the University of Aberdeen. Long-time participants and collaborators of the project (M. Gerken, P. Ebert, E. Zardini, A. McGlynn, S. Roca, D. Dodd and M. Smith) gave responses to the invited speakers and delegates from all over Europe, North America and Australia were in attendance. The event took place in a very enjoyable intellectual and social atmosphere. As detailed below, the conference registered exciting convergences in lines of research and overall provided a state-of-the-art showcase of research on apriority conducted by some of the leading figures in the field.

A. Casullo ("Challenging the A Priori/A Posteriori Distinction") offered a taxonomy of the attacks against the a priori/a posteriori distinction. As an instance of the kind of attack contending that the notion of apriority contains elements that are in tension with certain other features usually associated with it, he argued in favour of analysing a priori justification as justification based on a non-experiential source, and added that, on this analysis, the "default justification" variously postulated by some theorists (H. Field and C. Wright) would turn out to be neither a posteriori nor a priori.

P. Boghossian ("Intuition and the A Priori") discussed some problems with the understanding-based account of a priori justification. He then turned to the alternative intuition-based account, arguing that, in order to make sense of justification for inference, this requires the existence of non-propositional intuitions. He conjectured that some such intuitions (more specifically, intuitions of identity) should independently be postulated in order to make sense of the transparency of mental content.

B. Jarvis and J. Jenkins Ichikawa ("Apriority and the Objectivity of Rational Inquiry") defended a view according to which, in one sense of 'justification', everyone always has justification for a priori propositions, independently of whether one enjoys certain intuitions about them, and argued that it is this kind of justification that sets the standards for rationality.

M. Balcerak Jackson ("Imagination and the A Priori/A Posteriori Distinction") contended that imagination can be a source of a priori justification in the sense that, although its general workings are shaped by some very general features of our experiences, it does not serve to record specific facts about our surroundings.

I. Rumfitt ("Is Logic Empirical?") criticised H. Putnam's earlier views about the logic of quantum mechanics. He then offered a general anti-realist semantics that treats disjunction as a closure operation. Applying such semantics to the language of quantum mechanics, he showed that it invalidates the law of distribution.

C. Ichikawa Jenkins ("Justification Magnets") explored views according to which, although certain incompatible propositions are equally supported by the evidence, the tie is broken by a "justification magnet". In particular, she considered the proposals that justification gets attached to the alternative that is more natural, or to the alternative that is in fact true, discussing the pros and cons of both options.

D. Chalmers ("The Non-Modal Conception of Propositional Apriority") focussed on whether apriority for a proposition should be analysed in terms of the proposition's a priori knowability. He argued that such an analysis is committed to denying the necessary factivity of propositional apriority and that it is committed to rejecting the apriority of some instances of 'P iff actually P'. He developed an alternative analysis according to which propositional apriority consists in the existence of a certain kind of a priori propositional justification that is subject-independent.

L. BonJour ("In Defence of Rational Insight") argued that rational insight is grounded in non-propositional abilities of grasping properties, and offered a metaphysical sketch of how such abilities are possible. He also compared his theory with C. Ichikawa Jenkins' theory of concept examination, criticising her grounding requirement.

C. Wright ("The Basic A Priori: Arithmetic As a Case Study") critically explored several accounts of our perceptually aided basic arithmetical practices, including one according to which perceptual aids are tokens that are used for discovering new properties of their types. He argued that even this otherwise insightful account cannot explain how facts discovered about the tokens can be taken to reveal new facts about their types. He then proposed a Wittgensteinian non-cognitivist account on which the former facts enter as materials in the on-going creation of the type by the practitioners.

<div align="right">

ELIA ZARDINI
NIP, University of Aberdeen

</div>

Calls for Papers

INFORGS AND THE INFOSPHERE: THEMES FROM LUCIANO FLORIDI'S PHILOSOPHY OF ARTIFICIAL INTELLIGENCE: special issue of *The Journal of Experimental & Theoretical Artificial Intelligence*, deadline 1 July.

MIND AND PARADOX: special issue of *Journal of Experimental & Theoretical Artifical Intelligence*, deadline 1 July.

THE AIM OF BELIEF: special issue of *Teorema*, deadline 15 September.

SCIENCE VS. SOCIETY? SOCIAL EPISTEMOLOGY MEETS THE PHILOSOPHY OF THE HUMANITIES: special issue of *Foundations of Science*, deadline 31 October.

WHAT'S HOT IN . . .

. . . Logic and Rational Interaction

Despite the title of this column, this month we focus on interaction leading to irrational outcomes. To be a bit more precise: We focus on irrational beliefs held in states of *pluralistic ignorance*. This phenomenon, famously pictured in H.C. Andersen's tale of the emperors's new clothes, can be described as follows: An agent privately holds a certain conviction P, but observes the rest of society supporting the converse $\neg P$. Observing this public consensus brings him to utter $\neg P$ himself as his public opinion, even though he does not change his private belief. If this holds true of every agent we speak of pluralistic ignorance, a state in which the public opinion contradicts every individual opinion held.

Originally, the term was phrased in the 30ies, by Katz and Allport. After a long time where pluralistic ignorance was studied in social psychology exclusively,

the subject was recently discovered by various formal philosophers and logicians. It prominently featured as one of the four cases in the Copenhagen-Lund conference series last year.

To begin with, Bjerring, Hansen and Pedersen give a recent discussion of the topic, including the question whether pluralistic ignorance can occur as the result of rational interaction.

On the formal side, the recent months have seen several attempts to model the phenomenon. The attempts taken to understand pluralistic ignorance make use of plausibility logic with public announcement (Hansen) and dynamic doxastic logic (Olsson/Proietti).

On the other hand Lisciandra, Hartmann et al. work in a bayesian framework, using agent based simulations to study the creation and collapse of pluralistic ignorance.

A lot of interesting work on this topic can be expected to come up in the near future.

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 22–23 May 2012, LSE's Centre for the Analysis of Time Series (CATS) held a *Workshop on Uncertainty Quantification, Risk and Decision-making*. The programme was remarkably multi-disciplinary with talks covering as diverse topics as the accuracy of weather forecasts, carbon-related policy-making, engineering risk, economic and financial modelling, computational simulations and the foundations of statistical reasoning.

One interesting aspect that emerged from the talks and the panel discussions concerned the issue of communication, which in fact constitutes a three-fold problem in the wider area of uncertain reasoning and decision-making.

First, distinct research communities need to talk to each other. The workshop gave an interesting example of the fact that effective communication can take place even in the absence of a fully shared language. Albeit different people phrased it (very) differently, participants of this workshop appeared to agree on a fun-

damental distinction which came under various names including: risk vs uncertainty, aleatory uncertainty vs epistemic uncertainty, uncertainty vs ambiguity, uncertainty vs model-uncertainty and even ignorance vs randomness. Since it is unlikely that an obviously flawed approach can last for long in the scientific community, good reasons are to be expected to lie behind linguistic disagreement.

Second, the "scientific community" as a whole needs to inform decision-makers and make sure that the message gets across as uncorrupted as possible. This is particularly difficult in areas as complex as climate science, where there is hardly internal agreement within the scientific community in the first place, and even when this is possible, it seldom takes the form of a straightforward quantification. Models often disagree and even if we agree that the disagreement is measurable, decision-makers often just want the expert's opinions to be "right".

Third, scientists, experts and policy-makers need to talk to "the stake-holders", i.e., us. It goes without saying that in matters of uncertainty quantification and risk management, what might be a suitable language for one channel of communication need not be suitable for the other. So, whilst we might agree that a doctor's uncertainty is best understood as subjective, epistemic uncertainty, it might nonetheless be better expressed in a frequentist language (see, e.g., Spiegelhalter, D.J. (2011) 'Quantifying uncertainty' in *Risk*, edited by L Skinns, M. Scott, and T. Cox, 17–33, Cambridge Univestity Press).

To me, this rather daunting three-fold challenge suggests that we should aim at representing diversity in a framework which is uniform enough to guarantee effective communication. This might be achieved by putting a greater emphasis on a problem-based approach to the *foundations* of uncertain reasoning. Indeed, reasoning about the foundations can easily be supplied with practical problems, which in turn can usefully feed back on the foundations. Maybe this mechanism might provide the right amount of linguistic coordination that we need to communicate at all the above three levels.

In his talk, for instance, Jochen Bröcker, put forward an intuitive distinction between "weather forecasts" and "climate forecasts". Whilst weather forecasts, e.g., whether it will rain tomorrow at 12:00 in Pisa, tend to be extremely reliable, it is a lot harder to tell if Pisa "is getting drier".

Massimo Marinacci illustrated an approach for separating agent-dependent from agent-independent sources of uncertainty in decision problems and suggested its potential application in portfolio-optimisation problems.

Bernard Sinclair-Desgagné focussed on policy-making under model disagreement and put forward a proposal based on the willingness-to-pay to remedy to a currently unsatisfactory status quo.

Problems of the sort discussed in those three talks clearly call for criteria for model-construction and model-selection which are specific to the real-world problem at hand. A number of other talks presented case-studies about how uncertainty and risk are understood and measured in specific domains along with an overview of the challenges which remain open there. I think this workshop was very successful at raising mutual awareness among contiguous yet distinct areas in which the quantification of uncertainty plays such a fundamental role.

Hykel Hosni
Scuola Normale Superiore, Pisa
CPNSS, LSE

Events

July

Markets, Mechanisms, and Multi-Agent Models: ICML Workshop on the Interaction of Machine Learning and Economics, Edinburgh, 1 July.

MLG: 10th workshop on Mining and Learning with Graphs, Edinburgh, 1 July.

STAMLINS: ICML Workshop on Statistics, Machine Learning and Neuroscience, Edinburgh, 1 July.

AAP2012: Conference of the Australasian Association of Philosophy, University of Wollongong, 1–6 July.

Science and Metaphysics: University of Kent, 2–3 July.

Uncertainty in Computer Models: Sheffield, UK, 2–4 July.

AISB/IACAP: Birmingham, UK, 2–6 July.

HAI: Hypercomputation and AI Symposium, Birmingham, UK, 2–6 July.

LASR: 31st Leeds Annual Statistical Research Workshop, University of Leeds, 3–5 July.

Bounded Rationality: Summer Institute on Bounded Rationality, Berlin, Germany, 3–10 July.

Foundations for an Interdisciplinary Decision Theory: Max Planck Institute for Human Development, Berlin, Germany, 3–10 July.

ICT: 7th International Conference on Thinking, London, 4–6 July.

IIBM: 5th International Workshop on Intelligent Informatics in Biology and Medicine, Palermo, Italy, 4–6 July.

History and Philosophy of Programming: Ghent University, 5–6 July.

BSPS: Annual Conference of the British Society for the Philosophy of Science, University of Stirling, 5–6 July.

Joint Session: The Aristotelian Society and the Mind Association, University of Stirling, 6–8 July.

CAV: 24th International Conference on Computer Aided Verification, Berkeley, 7–13 July.

ISSCSS: International Summer School in Cognitive Sciences and Semantics, Latvia, 8–18 July.

ASC: 21st Australian Statistical Conference, Adelaide, 9–12 July.

RSC: 35th Annual Research Students' Conference in Probability and Statistics, University of Southampton, 9–12 July.

IPMU: 14th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Catania, Italy, 9–13 July.

ICALP: 39th International Colloquium on Automata, Languages and Programming, University of Warwick, 9–13 July.

WorldCong: 8th World Congress in Probability and Statistics, Istanbul, 9–14 July.

Foundations of Mathematics: University of Cambridge, 10–12 July.

TViTC: Theoretical Virtues in Theory-Choice, University of Konstanz, 12–14 July.

Logic Colloquium: University of Manchester, 12–18 July.

ICNCI: International Conference on Network and Computational Intelligence, Haikou, China, 14–15 July.

DEON: 11th International Conference on Deontic Logic in Computer Science, University of Bergen, Norway, 16–18 July.

WorldComp: The 2012 World Congress in Computer Science, Computer Engineering, and Applied Computing, Nevada, USA, 16–19 July.

DMIN: 8th International Conference on Data Mining, Nevada, USA, 16–19 July.

SIPTAss: Society for Imprecise Probability: Theories and Applications Summer School, Pescara, Italy, 16–20 July.

HUJI: Graduate Conference in Philosophy, Hebrew University of Jerusalem, 18–19 July.

Interfaces of the Mind: workshop at Ruhr-Universität Bochum, Germany, 19–21 July.

ASLP: Annual Conference of Australian Society of Legal Philosophy, Macquarie University in Sydney, 20–22 July.

ISA: IADIS International Conference Intelligent Systems and Agents, Lisbon, Portugal, 21–23 July.

ICL: 19th International Congress of Linguists, Geneva, Switzerland, 22–27 July.

Paradox and Logical Revision: LMU, Munich, 23–25 July.

WoMO: 6th International Workshop on Modular Ontologies, Graz, Austria, 24 July.

FOIS: 7th International Conference on Formal Ontologies in Information Systems, Graz, Austria, 24–27 July.

Einstein's Philosophy of Science: Summer School, University of Tübingen, 30 July–3 August.

August

NAFIPS: 31th North American Fuzzy Information Processing Society Annual Conference, Berkeley, 6–8 August.

CLAM: Logic and Computability Session, Latin American Congress of Mathematicians, Argentina, 6–10 August.

PMUV: Philosophy and Mathematics of Uncertainty and Vagueness, Brazil, 6–15 August.

ESSLLI: 24th European Summer School in Logic, Language and Information, Poland, 6–17 August.

KDD: 18th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Beijing, China, 12–16 August.

StaRAI: 2nd Statistical Relational AI workshop, Cataline Island, USA, 13 August.

ITP: 3rd Conference on Interactive Theorem Proving, Princeton, NJ, 13–16 August.

Logic and Cognition: Logic and Cognition Workshop, Opole, Poland, 13–17 August.

Historical Counterfactuals: Workshop, Bristol, 14 August.

UAI: Conference on Uncertainty in Artificial Intelligence, Catalina Island, USA, 15–17 August.

BMAw: 9th Bayesian Modeling Applications workshop, Catalina Island, 18 August.

UiNI: Uncertainty in Natural Intelligence workshop, Catalina Island, 18 August.

SLS: 8th Scandinavian Logic Symposium, Roskilde University, Denmark, 20–21 August.

ALFAn: Latin American Analytic Philosophy Conference, Buenos Aires, Argentina, 21–24 August.

CTF: Concept Types and Frames in Language, Cognition, and Science, Düsseldorf, 22–24 August.

AIML: Advances in Modal Logic, Copenhagen, 22–25 August.

FMIP: Munich / Groningen Summer School: Formal Methods in Philosophy, Groningen, 23–28 August.

EASLLC: International Conference and the Second East-Asian School on Logic, Language and Computation, Chongqing, China, 25–31 August.

FLINS: 10th International FLINS Conference on Uncertainty Modeling in Knowledge Engineering and Decision Making, 26–29 August.

CLIMA: 13th International Workshop on Computational Logic in Multi-Agent Systems, Montpellier, France, 27–28 August.

STeDy: Spatio-Temporal Dynamics workshop, Montpellier, France, 27–28 August.

WL4AI: Weighted Logics for AI workshop, Montpellier, France, 27–28 August.

AIGM: Algorithmic Issues for Inference in Graphical Models, Montpellier, France, 27–28 August.

ARCOE: 4th International Workshop on Acquisition, Representation and Reasoning with Contextualized Knowledge, Montpellier, France, 27–28 August.

ECAI: 20th European Conference on Artificial Intelligence, Montpellier, France, 27–31 August.

COMPSTAT: 20th International Conference on Computational Statistics, Cyprus, 27–31 August.

Collective Intentionality: University of Manchester, 28–31 August.

CNL: Workshop on Controlled Natural Language, Zurich, 29–31 August.

FoR&D: Conference on Frontiers of Rationality and Decision, University of Groningen, 29–31 August.

September

CSL: 21st EACSL Annual Conference on Computer Science Logic, Fontainebleau, France, 3–6 September.

WoLLIC: Workshop on Logic, Language, Information and Computation, Argentina, 3–6 September.

ABS: Applied Bayesian Statistics School, Italy, 3–7 September.

ICLP: 28th International Conference on Logic Programming, Budapest, 4–8 September.

ıKNOW12: 12th International Conference on Knowledge Management and Knowledge Technologies, Graz, Austria, 5–7 September.

ECıtS

Evidence and Causality in the Sciences, University of Kent, 5–7 September

---

GAMES: Games for Design and Verification, Napoli, Italy, 7–12 September.

Intuitions, Experiments and Philosophy: University of Nottingham, 8–9 September.

Logic and Relativity: 1st International Conference on Logic and Relativity, Budapest, 8–12 September.

WEO-DIA: 1st Workshop on Well-founded Everyday Ontologies–Design, Implementations & Applications, Wroclaw, Poland, 9 September.

COMMA 2012: 4th International Conference on Computational Models of Argument, Vienna, Austria, 10–12 September.

LATD: Logic, Algebra and Truth Degrees, Japan, 10–14 September.

WPMSIIP: 5th Workshop on Principles and Methods of Statistical Inference with Interval Probability, Munich, Germany, 10–15 September.

Datalog 2.0: 2nd Workshop on the Resurgence of Datalog in Academia and Industry, Vienna, Austria, 11–14 September.

L&R: workshop on Lattices and Relations, ILLC, University of Amsterdam, 12–14 September.

WUPES: Workshop on Uncertainty Processing, Czech Rep., 12–15 September.

ENFA: 5th Meeting of the Portuguese Society for Analytic Philosophy, University of Minho, Braga, 13–15 September.

SOPhiA: Salzburg Conference for Young Analytic Philosophy, University of Salzburg, Austria, 13–15 September.

Colloquium Logicum: Paderborn, Germany, 13–15 September.

SIFA: 10th National Conference of the Italian Society for Analytic Philosophy, Alghero, 13–15 September.

SUM: 6th International Conference on Scalable Uncertainty Management, Marburg, Germany, 17–19 September.

ILP: 22nd International Conference on Inductive Logic Programming, Dubrovnik, 17–19 September.

GAP8: 8th Conference of the Society for Analytic Philosophy, Germany, 17–20 September.

Logical Form: University of Cambridge, 18–19 September.

SemDial: 16th Workshop on the Semantics and Pragmatics of Dialogue, Université Paris-Diderot, 19–21 September.

PGM: 6th European Workshop on Probabilistic Graphical Models, Granada, Spain, 19–21 September.

FORMAL METHODS IN ARGUMENT RECONSTRUCTION: Konstanz, Germany, 20–21 September.

CALINTSS: Causation and Laws in the Special Science—Metaphysical Foundations, Konstanz, 21–22 September.

PHILOSOPHICAL ISSUES IN BELIEF REVISION, CONDITIONAL LOGIC AND POSSIBLE WORLD SEMANTICS: Konstanz, Germany, 21–22 September.

ENPOSS: 1st European Network for the Philosophy of the Social Sciences Conference, University of Copenhagen, 21–23 September.

MLSP: IEEE Workshop on Machine Learning for Signal Processing, special session on Causal Discovery, Spain, 23–26 September.

STRUCTURE AND UNCERTAINTY: workshop on Modelling, Inference and Computation in Complex Stochastic Systems, Bristol, 24–27 September.

TISS: Tübingen International Summer School—How do we make decisions?, 24–27 September.

ECML-PKDD: European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Bristol, UK, 24–28 September.

JELIA: 12th European Conference on Logics in Artificial Intelligence, Toulouse, 26–28 September.

CONSCIOUSNESS AND VOLITION: 1st International Krakow Conference in Cognitive Science, Krakow, Poland, 27–29 September.

LNK: 5th Conference on Non-Classical Logic. Theory and Applications, Poland, 27–29 September.

MEW6: 6th annual Midwest Epistemology Workshop, Indiana University, Bloomington, 28–29 September.

LSFA: 7th Workshop on Logical and Semantic Frameworks with Applications, Rio de Janeiro, 29–30 September.

OCTOBER

DEPARTING FROM SAINSBURY: University of Barcelona, 1–2 October.

SMPS: 6th International Conference on Soft Methods in Probability and Statistics, Konstanz, 4–6 October.

FPMW: 4th French PhilMath Workshop, Collège de France, Paris, 4–6 October.

PHILOSOPHY OF SCIENTIFIC EXPERIMENTATION: University of Colorado, Boulder, 5–6 October.

TIC2: Turing in Context II: Historical and Contemporary Research in Logic, Computing Machinery and AI, Brussels, 10–12 October.

FORMAL ETHICS: Munich, 11–13 October.

THE ROLES OF EXPERIENCE IN A PRIORI KNOWLEDGE: University of Cologne, Germany, 13–14 October.

NUMBERS & TRUTH: The Philosophy and Mathematics of Arithmetic and Truth, University of Gothenburg, Sweden, 19–21 October.

ATAI: Advanced Topics in Artificial Intelligence, Bali, Indonesia, 22–23 October.

ECREA: 4th European Communication Conference, Istanbul, Turkey, 24–27 October.

IDA: 11th International Symposium on Intelligent Data Analysis, Helsinki, Finland, 25–27 October.

COURSES AND PROGRAMMES

Courses

NASSLLI: North American Summer School in Logic, Language and Information, University of Texas at Austin, 18–22 June.

EINSTEIN'S PHILOSOPHY OF SCIENCE: Summer School, University of Tübingen, 30 July–3 August.

ESSLLI: 24th European Summer School in Logic, Language and Information, Opole, Poland, 6–17 August.

FMIP: Munich / Groningen Summer School: Formal Methods in Philosophy, Groningen, 23–28 August.

TISS: Tübingen International Summer School—How do we make decisions?, 24–27 September.

Programmes

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

Post-doc position: in Logic and / or Philosophy of Mathematics, Université du Québec à Montréal, until filled.

Lecturer: in Statistics, University of Manchester, until filled.

Associate Professor of 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: on the project "Explanatory Reasoning: Normative and Empirical Considerations," Tilburg Center for Logic and Philosophy of Science, 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.

Teaching Fellow: in Logic and History of Analytic Philosophy, University of Edinburgh, deadline 10 July.

Postdoc Position: on EPSRC-funded project "Classical Dependent Type Theories and Classical Logic-Enriched Type Theories," Royal Holloway, University of London, deadline 16 July.

Three Assistant Professorships: Munich Center for Mathematical Philosophy, LMU, deadline 18 July.

Six Postdoc Positions: Munich Center for Mathematical Philosophy, LMU, deadline 18 July.

Postdoc Position: on the project "Explaining Language: Philosophical Perspectives on Computational Linguistics," TiLPS, Tilburg University, deadline 15 August.

Studentships

Two Doctoral Training Grants: School of Computing, Faculty of Engineering, University of Leeds, until filled.

PhD position: in Bayesian Decision Theory, School of Computer Science and Statistics, Trinity College Dublin, until filled.

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

PhD positions: in Statistical Methodology and its Application, University College London, until filled.

PhD position: in Logic and Theoretical Philosophy at the Institute for Logic, Language and Computation at the University of Amsterdam, until filled.

Six PhD positions: Munich Center for Mathematical Philosophy, LMU, deadline 18 July.

Four PhD positions: in "Foundations of the Life Sciences and their Ethical Consequences," European School of Molecular Medicine, University of Milan, deadline 3 September.

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.

University of **Kent** CENTRE FOR REASONING

www.kent.ac.uk/reasoning