The Reasoner

Volume 8, Number 10
October 2014

www.thereasoner.org
ISSN 1757-0522

Contents

Editorial

Much as I love philosophy, I have worried in the past that it is just a self-indulgent activity. Shouldn't we be actually making a difference, somehow, or contributing to

some wider debate? I've since come to see that work in philosophy can be relevant to all sorts of issues. As a case study, this issue features the Managing Severe Uncertainty Project at the LSE's department of Philosophy, Logic and Scientific Method.

The project was set up in March 2013, and there are nine members: Richard Bradley, Roman Frigg, Katie Steele, Alex Voorhoeve, Charlotte Werndl, Hykel Hosni, Casey Helgeson, Thomas Rowe and Silvia Milano. The project has a focus on climate change—an area on which policy decisions must be made despite severe and multi-faceted uncertainty. The project members meet frequently, and I have attended some of these meetings and seen them in action. Early in the project, the meetings took the form of 'masterclasses' where the members pooled their knowledge. Since then, there have been meetings focused on members' current research in the area, sessions spent analysing the reports produced by the IPCC (Intergovernmental Panel on Climate Change) and papers by Nicholas Stern, and planning meetings to discuss the project's future direction of research. The group has invited in external speakers and advisors from the economics department, the Grantham Research Institute on Climate Change and the Environment, and the Centre for the Analysis of Timeseries, and in turn members of the project have disseminated their work to a wide range of audiences internationally.

I felt that the best way to introduce you to the project would be to give you a snapshot of the current work of some of its members. I hope you enjoy reading my (brief!) interviews with these five philosophers, and seeing where their interests converge.

Anna Mahtani
Philosophy, LSE

Features

Interview with the Managing Severe Uncertainty Group

Anna Mahtani: Richard, can you tell me how the project came about?

Richard Bradley: The origins of the project really lie in encounters between myself and Roman Frigg. Roman has a long-standing interest in climate change, and he had come to realise that it wasn't enough to complain about the limitations of the climate models that were being used: he also wanted to address the question of what kinds of outputs would be useful for policy makers. At the same time I was starting to think about how we make decisions in situations where we don't have precise probabilistic information. These two issues are obviously closely related. Roman was already working with Charlotte Werndl on these issues, and I was already talking actively to Katie Steele, so there was this very natural confluence of interest in this question.

I think this is what is special about the project: it's not just a philosophy of science project about modelling, it's not just decision theory, but it is specifically trying to look at that territory that falls between the two, which is very under-explored.

AM: What is severe uncertainty?

RB: Well, at first I was thinking of situations where you are unable to form precise probabilities for the relevant events. It has since become apparent to me that the expression 'severe uncertainty' isn't used uniquely in that way. For example, the expression is also used to describe situations where one might not be aware of the entire state space. And we can also use the expression to describe cases where there is some deep ethical uncertainty: these are cases where it's not that you don't know what the consequences of your actions are, it's just that you don't know what significance to attach to those consequences—whether they are good or bad. In some cases individuals or countries might have strong views on the moral significance of an action, but it's hard to tell on aggregate whether the action would be good or bad: here the individual level disagreement gives rise to group level uncertainty.

Richard Bradley

AM: The group is focusing its attention on climate change. Is that an issue that you care personally about?

RB: It's one of the most important problems that humankind faces at the moment. But it's not obvious to me that this is a battle between defenders of the environment and polluters of the environment. The problem is in a sense too serious for that. There's a collective fate here if we don't get things right: all of our future generations will be very severely affected by this.

AM: What would be the ideal outcome from the project?

RB: Well, the title of the project is *Managing* severe uncertainty, rather than *Solving* severe uncertainty. Uncertainty isn't going to go away—there is no miracle way of turning lack of evidence and disagreement into consensus and good evidential support. But we do feel that there are better and worse ways of handling these difficult situations, and one thing that we can do is acknowledge that there are trade-offs to be made here. Broadly, if we are willing to accept less precision, then we can get more certainty—for example we can either be relatively certain that the temperature will fall within some broad range, or less certain that it will fall within some narrower range. There should be a two-way communication between climate scientists and policy makers to establish what sort of information would be useful for decision making—and part of the aim of the project is to create a framework in which that is understood to be important.

AM: So Roman, you recently gave a keynote talk at the Understanding Risk Forum, organized by the World Bank. Was Boris Johnson a hard act to follow?

Roman Frigg: Not at all—he didn't come and sent his deputy instead! But that didn't make it less daunting. Giving a public lecture is a difficult task and I am always nervous when I have to give one.

AM: In your talk at this forum, you argued against more money being spent on climate modelling. Can you explain why?

RF: This statement has to be qualified. I didn't argue against climate modelling *tout court*—climate models are great

tools for many purposes. What I argued against is the use of climate models for the purpose of generating *local* projections. The *United Kingdom Climate Impacts Programme's* UKCP09 project makes high-resolution projections of the climate out to 2100: probabilities are given for events on a 25km grid, which means, for instance, that the projections may differentiate between the impacts of global climate change in London and Oxford. Given the acknowledged systematic errors in all current climate models, treating model outputs as decision relevant projections can be seriously misleading. In extrapolatory situations, such as projections of future climate change impacts, there is little reason to expect that post-processing of model outputs can correct for the consequences of such errors. This casts doubt on our ability, today, to make trustworthy, high-resolution probabilistic projections out to the end of this century. This, however, is not to say that climate models cannot be useful instruments in many ways. They can be used, for instance, to gain an understanding of various processes in the climate systems, and they can inform us about the effects of climate change on a global and continental scale.

Roman Frigg

AM: If we can't use the models to inform our decisions, then how do you think we should be making decisions on whether and how to mitigate the effects of climate change?

RF: It depends what decisions we have in mind. There is now a widespread consensus that global warming is real and in large part due to human activities. This is enough for mitigation: we know that we have to reduce greenhouse gas emissions. Things get more involved when we try to devise detailed adaptation strategies. The impact of climate change on humans (as well as other organisms) occurs at a local scale, and so ideally one would like to know what changes one has to expect in one's immediate environment. As I have argued above, projects like UKCP09 don't fit the bill. This leaves us the difficult question of what else would. I have no definitive answer to this question. An idea I would like to pursue is whether structured expert elicitation leads to better results. This approach has been used successfully in other contexts and it seems promising in the climate case too. What we know for sure is that we need to change our attitude to the problem. Scientists have come around to the idea that the aim of research is to understand, classify and communicate uncertainty rather then to reduce it, and policy-makers have to renounce the still popular first-predict-and-then-act principle and start accepting that climate decisions have to be made under uncertainty.

AM: Charlotte, you're interested in the question of how to define 'climate'. How is it normally defined?

Charlotte Werndl: The first thing to say is that there is no 'normal' definition—it depends which sub-community you are talking to.

Many climate scientists will tell you that climate is 'an ensemble distribution'. And this is a very weird idea—it's not intuitive at all. To take a simple example, suppose that we're just interested in temperature—say the temperature a year from today. Today's prediction of that temperature is a *distribution*

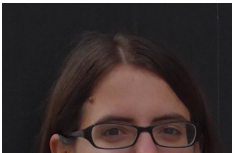

of possible temperatures—namely the distribution arising by evolving forward our uncertainty about today's temperature—and on this view, that distribution is the climate! Philosophers immediately think this idea is very strange, because on this view facts about climate are dependent on our uncertainty.

Then there's another sort of definition of 'climate', which is a distribution over time. For example, if you are interested in just temperature, then we might have a graph charting the temperature month by month. This is the sort of thing that you would call the 'climate' of a country when you learn geography at school. The usual idea is that we chart this distribution given constant external conditions—for example, we would hold constant the amount of energy coming from the sun.

You can imagine how having these two different definitions causes problems. For example when a climate modeller (who is working with the 'ensemble distribution' definition) talks to the policy advisor (who has something like the distribution over time definition in mind), there is a danger that they may be talking past each other.

AM: Which definition of 'climate' do you prefer?

CW: I don't like the ensemble distribution definition because it makes climate dependent on uncertainty. I like the distribution over time, but I don't like the requirement that the external conditions should be held constant. On my view, climate is a distribution over time under a *regime* of varying external conditions. For this new definition of climate, we need a new mathematics, and that is something that I am working on. This is a very young field of mathematics: people have just started work on this in the past ten years, so it's an exciting area.

AM: How should we understand climate *change* under these different definitions?

CW: Well, some definitions of climate can't make sense of climate change. The distribution over time definition can: for example, take the regime from 1st July 1970 to 1st July 2000, and take the regime from 1st July 2000 to 1st July 2030: we can compare the two distributions of temperature over time—and see if they are different.

For the ensemble distribution, though, it's much harder to see what talk about climate change means. Think about this: what was the climate in the past—say in 1950? Here we don't have uncertainty, so what is the climate? Climate scientists usually don't know how to answer this question, but one answer I've been given is this: maybe the climate in 1950 is the distribution that scientists in 1900 could have predicted! But this is arbitrary—why 1900 rather than 1930, say? And of course climate change depends on whether past climate is different from present or future climate, so I don't think that with the ensemble definition we can make sense of climate change.

AM: If people can't even agree about the correct definition of climate change, then is there a consensus that climate change is actually happening?

CW: Well, when we worry about climate change, we worry about something specific, namely that the average temperature is going up—and there is no doubt about that. This point is independent of how we define 'climate'. To some extent the question of how to define 'climate' is a linguistic question, but it's an important one—because we need to be able to communicate. When the policy advisor is talking to the climate

scientist, for example, they need to know what they both mean by the 'temperature distribution'.

AM: Alex, you've been doing a lot of work in the LSE economics lab. But what is an economics lab?

Alex Voorhoeve: The LSE Behavioural Economics lab is a set of networked computers which can be programmed with questionnaires, interactive games, and gambles. It recruits subjects from across London, who enrol in these experiments for pay.

AM: What do you think your experiments have shown?

AV: Together with an economist, Ken Binmore, a psychologist, Lisa Stewart, and a statistician, Arnaldur Stefansson, I've done experiments on how our decision-making under ambiguity (when we don't have the probabilities of possible outcomes) differs from decision-making under risk (when we do). In our variations on the classic Ellsberg experiments, subjects had to choose between a risky gamble and an ambiguous gamble (where they knew only the possible prizes, but not their probabilities). A common finding in such cases is ambiguity aversion—people prefer gambles with known probabilities, in a manner that is inconsistent with Bayesian decision theory, which requires that they choose as if they assign probabilities to the ambiguous events. However, to our surprise, we found very little ambiguity aversion. Instead, the best predictor of subjects' behaviour was the classic "principle of insufficient reason". This refuted our own theory of decision under ambiguity and the theories of others (Popper would have been proud). Intrigued, we varied our experiments. The results are just in, and we have found that supposedly important changes in 'framing' (such as whether gambles involve gains or losses) make very little difference. Moreover, individuals who follow a coherent strategy appear to use Bayesian decision theory. Beyond the conclusion that commonly reported results are suspect, I hypothesize that the famous Ellsberg set-up does not involve enough ambiguity. It is too tame, too manageable. To *really* test ambiguity aversion, we have to expose people to choices in which they don't even know what all the possible outcomes are.

Alex Voorhoeve

AM: For a real-life case like climate change—where policy makers need to act with limited information—what do your results suggest? Can you use them to predict what sorts of decisions policy makers will take? Or to judge what decisions *should* be made?

AV: Theories that predict (or prescribe) ambiguity aversion have been used in the evaluation of climate policies—e.g., in arguments about the size of an optimal carbon tax. Insofar as these theories gained support from the standard findings in Ellsberg-type cases, our experiments cast doubt on them. But that doesn't mean I would favour using standard Bayesian decision theory! Rather, climate change seems to me to involve much more severe ambiguity than these theories allow, since we don't even know what all the possible outcomes might be, or how to evaluate them. So I'm afraid that neither standard decision theory, nor familiar departures from it are as yet in a position to predict or prescribe how we should decide.

AM: Casey, you joined the LSE specifically to work on this project. What attracted you to it?

Casey Helgeson: My most rewarding research experiences have been collaborative, so the idea of being part of a team was very appealing. Also, I like working at the intersection of fields. The possibility of some positive impact beyond academia is also a motivation for me personally—not that more "pure" philosophy isn't also valuable.

AM: I understand that you have been reading a lot of documents produced by the IPCC. What are these documents?

CH: IPCC (Intergovernmental Panel on Climate Change) reports aim to summarize the current state of knowledge about the changing climate, its impacts, and the prospects for mitigation and adaptation. I'm reading bits of the most recent round of reports that came out earlier this year (AR5), and summaries of those. A lot of it is painful reading. The Summaries for Policymakers can get jumbled and muddied in the final, very politicized editing process (see John Broome's London Review of Books blogpost). The chapters and technical summaries are scientifically better, but their encyclopaedic structure and careful language puts me right to sleep. You have to work through it though, because the content is really interesting and important, and it's a great service to have so much research summarized in a timely way. It's a massive effort to put these reports together.

Casey Helgeson

AM: What are you looking for in these documents?

CH: Mainly, I'm looking at how uncertainty is expressed. Presenting the state of scientific knowledge of course requires addressing uncertainty, and this has gradually become more and more central to the reports. Now there's a guide for IPCC authors on how to express uncertainty. And their spiel on uncertainty, and their framework for expressing it, is front and center in everything they produce. They use a two-level approach where a statement can get a probability (or range of probabilities), but then it also gets a level of "confidence" that reflects the depth of scientific understanding behind the statement (and if applicable, the probability). The two can stack, so you end up with statements like "Such and such change in precipitation is *very likely* (high confidence)."

Results expressed in this framework have several purposes, one of them being to convey information for decision-making. So then questions arise, like how the two aspects of uncertainty should enter into good decisions. Just to quickly indicate that there is a real question here, standard expected utility maximization (for example) has no room for something like "confidence." So do you throw it away? Do you collapse likelihood and confidence into a probability? Or do you expand your approach to decision making to make room for both?

AM: Are you going to take your results to the IPCC?

CH: Well, it's funny to think about it like that, because it's not like there's anyone in particular to tell. There's a community of scientists who produce research on climate change, and then there's a (surprisingly large) subset of those folks who periodically summarize everything that's out there to create these reports within the IPCC frame-

work. If one does good research on any relevant aspect of climate change—and the uncertainty framework is now a part of that—by default that becomes a part of what the IPCC summarizes (assuming there will be a sixth assessment report).

But there are also bigger-picture reasons to study the IPCC's uncertainty framework. Research on how to represent deep uncertainty is relevant to other sciences as well, and to decision-making in finance, security, public health, etc. The urgency and controversy surrounding climate change has pushed the IPCC to develop a very explicit approach to communicating uncertainty; lots of scientists are using it, lots of policymakers are seeing it; their framework has the potential to become a paradigm in the communication of scientific uncertainty. So we'd better figure out how best to use information presented like that—or, if there is something better, change the framework itself before it sticks.

Transitivity and High Probability

Peter Achinstein (2001: *The Book of Evidence*, Oxford: Ch. 6) argues that evidence must satisfy a Condition of High Probability **HP**:

> **HP** If e is evidence for h, or, alternatively put, if e is a good reason to believe h, then $\mathrm{P}(h|e) > .5$.

The intuitive argument for **HP** is that if e were a good reason to believe h, but $P(h|e)$ were nevertheless less than or equal to .5, then e would seem to be at least as good a reason, if not a better reason, to believe $\neg h$. But, Achinstein assumes, a good reason to believe h cannot also be a good reason to believe $\neg h$.

John Pollock (1991: "A Theory of Defeasible Reasoning," *International Journal of Intelligent Systems*, 33–54), accepts a version of **HP** as well. He defines the "strength" of a reason as $\log(.5/1 - r)$, where r is the value of one of a standard set of conditional probabilities he uses as a benchmark. On this measure, if r is less than .5, the strength of the reason is 0.

A puzzle emerges, however, in that Pollock couples his acceptance of **HP** with the further thesis that good reasons are transitive. Pollock (1991: 47–48) thinks transitivity must hold because "[o]nce one has arrived at a set of conclusions, one does not hesitate" to use those conclusions in further inferences. (Pollock (1990: *Nomic Probability and the Foundations of Induction*, Oxford: Chs. 3, 8) provides additional discussion.) Pollock thus accepts:

> **T** If R is a good reason to believe S and S is a good reason to believe T, then R is a good reason to believe T.

Putting **HP** and **T** together yields a condition on chains of good reasons that might be called Transitivity of High Probability **THP**:

> **THP** If R is a good reason to believe S and S is a good reason to believe T, then $P(T|R) > .5$.

It is well known that various Bayesian measures of probabilistic support are intransitive. (Atkinson and Peijnenburg (2013: "Transitivity and Partial Screening Off," *Theoria*, 294–308) provide a good overview of the literature.) The failure of transitivity for these measures should make us question **THP** as well.

To keep the discussion simple, suppose that we restrict ourselves to cases where the prior or unconditional probabilities of R, S and T are equal and exceed 0, for example, where they all equal .5. Suppose further that, consistent with **HP**, both $P(S|R)$ and $P(T|S)$ equal or exceed some number k greater than .5. Now,

(1) $P(\neg T|R) = P(\neg T\&S|R) + P(\neg T\&\neg S|R)$.

Moreover,

(2) $P(\neg T\&\neg S|R) \leq 1 - k$,

and,

(3) $P(\neg T\&S|R) = P(\neg T\&R|S) \leq 1 - k$,

given that, by hypothesis, $P(S) = P(R)$. It follows that

(4) $P(\neg T|R) \leq 2 - 2k$,

and

(5) $P(T|R) \geq 1 - (2 - 2k) = 2k - 1$.

Thus, if, say, both $P(S|R) \geq .6$ and $P(T|S) \geq .6$, then $P(T|R) \geq .2$. While this is compatible with $P(T|R) > .5$, as **THP** requires, some additional substantive assumption will be needed, beyond **HP** and **T**, to secure this outcome.

The result (5) can be extended to chains of propositions of arbitrary length:

(6) $P(A_n|A_1) \geq 1 - (2^{n-2} - 2^{n-2}k) = 2^{n-2}k - (2^{n-2} - 1)$,

where, again, the unconditional probabilities of A_1 through A_n are equal and non-zero, and the conditional probabilities $P(A_i|A_{i-1})$ in the chain equal or exceed some number k. In such a chain,

(8) $P(A_n|A_1) > .5$ if $k > (2^{n-1} - 1)/2^{n-1}$.

Thus, where $n = 3$, $P(A_n|A_1) > .5$ if k exceeds $3/4$; where $n = 4$, $P(A_n|A_1) > .5$ if k exceeds $7/8$, and so on. By the same token, it is obvious that no matter how high we fix the value of k, there is an n such that the lower bound value of $P(A_n|A_1)$ drops below .5. Hence, it seems that **THP** must fail for some chain-length n, unless, again, some additional substantive principle beyond **T** and **HP** ensures that $P(A_n|A_1)$ always remains above .5.

One such additional substantive principle would be that R is a good reason to believe S only if $P(S|R) = 1$. This would be the case if, for example, R's being a good reason to

believe S required that R logically entail S. While this would secure **THP**, it is doubtless too stringent a requirement to impose on good reasons.

Another option would be to hold that whether R is a good reason to believe S somehow depends on the length of the chain of reasons preceding R. Put differently, on this suggestion, given a chain of reasons of length n such that $P(A_n|A_1) > .5$, the proposition A_n is a good reason to believe some further proposition A_{n+1} only if all of the conditional probabilities in the chain, including the newly-added link $P(A_{n+1}|A_n)$, are above a threshold k sufficient to ensure that $P(A_{n+1}|A_1) > .5$. This suggestion seems at best ad hoc. It is difficult to see how R's being a good reason to believe S could depend on where R and S are located in a larger chain of reasons, or how R could fail to be a good reason to believe S simply because R and S occur at the end of a long chain, rather than, say, at the beginning.

Yet another option would be to stipulate that A_1 is a good reason to believe A_n if none of the conditional probabilities $P(A_i|A_{i-1})$ in the chain is less than or equal to .5. This is essentially the solution that Pollock (1991: 49) adopts. Under his "Weakest Link Principle," the strength of support for the conclusion of a chain of reasons is the strength of the weakest reason in the chain. But ignoring the strength of other links in the chain seems dubious, and Pollock offers no motivation for it beyond the desire to preserve transitivity.

Short of adopting one of these seemingly unacceptable solutions, it appears that **T** or **HP** is wrong. Between the two, **T** seems the better candidate for rejection. This suggests in turn that Pollock is wrong in supposing that we can unhesitatingly use a conclusion in further inferences without looking back to the original reasons for that conclusion to determine whether they also support the additional inferences based on it.

Stephen A. Fogdall
Schnader Harrison Segal & Lewis LLP

News

The Evidence Workshop, 4–5 June

It is usually thought that a rational agent should proportion her beliefs to her evidence—only then will her beliefs be rational, justified, or qualify as knowledge. But just what counts as evidence is often the subject of much controversy. The Evidence Workshop was held at the Philosophy Department of the University of Kent, and brought leading philosophers together to discuss just what counts as evidence, with special attention given to Timothy Williamson's theory of evidence as knowledge.

There were approximately 25 participants, including philosophy and law graduate students from the UK, US, Poland, and Norway. The first day began with Paulina Sliwa (Cambridge) and her talk "Respecting all the evidence", in which she proposed a principle of evidential calibration in order to resolve problems caused by accommodating both first-order and higher-order evidence. This was followed by a response from Ruth Hibbert (Kent). Clayton Littlejohn (King's College, London) then spoke about "Epistemological disjunctivism and the basis of the basis problem" with a response given by

Michael Wilde (Kent). Littlejohn argued that the basis of the basis problem is an evidentialist account of the relationship between reasons for belief and appropriate epistemic status. The first day ended with the conference dinner, which was attended by all speakers and a large number of the participants.

The second day began with Michael Wilde (Kent) and his talk "Evidence and pragmatics" to which Jessica Brown (Arché, St Andrews) responded. Then came a mini-symposium of Timothy Williamson's theory of evidence as knowledge. Jessica Brown offered some criticisms of Timothy Williamson's infallibilism, to which Timothy Williamson (Oxford) gave a reply. Aidan McGlynn (Edinburgh) then argued that there are some gaps in Timothy Williamson's argument for his evidence as knowledge theory, and again Timothy Williamson gave a reply. Last, Jon Williamson (Kent) argued for an alternative to the knowledge theory of evidence, viz., evidence is that which a rational agent takes for granted in a given context. Timothy Williamson then concluded with some objections to Jon Williamson's proposal.

The workshop provided an occasion for focused and lively debate on an important and controversial topic in epistemology. It was organized by Julien Murzi (Kent / MCMP, Munich) and Michael Wilde (Kent), and hosted by the Kent Centre for Reasoning and the Southern Normativity Group. It was funded by the Centre for Reasoning, the Kent Institute for Advanced Studies in the Humanities, the School of European Culture and Languages, and the Mind Association.

Michael Wilde
Philosophy, University of Kent

Reasoning Club, 23–24 June

The third annual Reasoning Club conference took place at the University of Kent, and followed on from successful and enjoyable conferences at Brussels in September 2012 and Pisa in June 2013. The 2014 conference saw thirteen speakers contribute talks over a busy two days in June. The theme of the conference, as in previous years, was "Reasoning", broadly construed. The diversity of topics and points of view on offer led to enthusiastic and productive discussions following many of the talks.

The conference began with Katie Steele, the first of two keynote speakers, giving a talk on "Model-Selection Methods & the 'Use-Novelty' Criterion". Steele discussed Worrall's 'use-novelty (UN)' criterion, which says that fit-with-the-evidence should bear on our confidence in a theory when the evidence is new, in the sense that it was not accommodated in the construction of the theory. She argued that Worrall's theory is powerful because it makes the UN criterion explicable in terms of a Bayesian model of confirmation, but that Worrall's revised UN criterion deals only with a special case. Addressing the question of which model-selection methods make for plausible extensions of Worrall's UN criterion to more complicated cases, she contrasted the 'cross validation' method with the more traditional Bayesian account of model selection.

In the next session, Michael Wilde proposed a non-reductionist epistemic theory of causality, based on Timothy Williamson's conception of evidence, and Gil Sagi discussed a scale inspired by Ruth Barcan Marcus of various levels of meanings: exten-

sions, intensions and hyperintensions, proposing an extension of this scale to accommodate logical notions.

In the afternoon, Rafal Urbaniak analysed arguments that a Dutch Bookie is necessarily irrational, and Marta Sznajder gave a philosophical analysis of the presuppositions behind Rudolf Carnap's Gamma and Eta rules, and some problems they face. Then Jacopo Amidei put forward a qualitative foundation for non-standard probabilities, pointing out and justifying the representation of the binary relation "less probable than" by means of hyper-real probability functions, and Seamus Bradley defended the 'Imprecise Probabilities' model of belief against problems.

The second day began with our second keynote speaker, Richard Pettigrew, giving a talk on 'Reasoning and the accuracy-first approach to Epistemology'. Pettigrew's paper explored a problem with the Bayesian updating rule Conditionalization, which he argued arises due to misidentifying the role of Conditionalization. Rather than a rule that an agent must follow when she needs to accommodate new evidence, it is an inference rule that she must use when she is conducting probabilistic reasoning. He argued that an accuracy-first epistemology supports this conclusion.

In the next session, Thomas Schindler proposed a formulation of what it means for a theory of truth to enable one to express generalizations, and examined existing truth theories accordingly, and Joachim Frans, presenting joint work with Bart Van Kerkhove, proposed seeing explanation as a product of a reasoning process, and discussed the role of visualizations in mathematical explanations in the light of this.

In the afternoon, Rossella Marrano talked about the robustness of the truth-values versus belief-values distinction in Łukasiewicz's real-valued logic, and Lavinia Picollo assessed claims that the notion of validity understood as logical truth engenders contradictions. Finally, Jürgen Landes closed the conference with a talk explaining how to justify the principle of indifference and maximum entropy principles using scoring rules.

The conference was organised by Julien Murzi, Ruth Hibbert and Graeme Forbes, and was funded by the Kent Institute for Advanced Studies in the Humanities (KIASH), the School of European Culture and Languages (SECL) at the University of Kent, the British Society for the Philosophy of Science (BSPS), the British Logic Colloquium (BLC), and the Centre for Reasoning at Kent.

Ruth Hibbert
Philosophy, University of Kent

Cognitive Science of Science, 18–22 August

Polish Centre of Philosophical Research hosted a workshop on Cognitive Science of Science on August 18–22, in Kazimierz Dolny, Poland. The main goal was to bring together relevant work from a number of perspectives (including developmental cognitive psychology, evolutionary explanations of human behaviour, cognitive science, philosophy of science and history of science) and focus on the social, cultural and material underpinnings of the sophisticated cognition that science involves.

The workshop was opened by one of the keynote speakers—Cristine Legare (Uni-

versity of Texas at Austin, USA) who discussed the ontogeny of cultural learning. Dr Legare discussed her works on cognitive developmental account of how children flexibly use imitation and innovation as dual engines of cultural learning. Next, Yuichi Amitani (Tokyo University of Agriculture, Japan) presented his talk on the application of psychology to resolve the problem of biological species without definitions. John Collier (University of KwaZulu-Natal, South Africa) discussed the strange case of female Hormone Replacement Therapy and Cognitive Salience. He argued that salience plays a major role in the acceptance or rejection of postmenopausal testosterone replacement for women, especially the latter. Next, Corinne L. Bloch-Mullins (Marquette University, USA) examined the role of mental categorization in investigative practice, apart from the integration of existing information or the application of the theoretical knowledge that may be represented by the concept. Professor Bloch-Mullins argued that concepts are, in a very strong sense, tools of discovery. Miles MacLeod (University of Helsinki, Finland) provided a case study from integrative systems biology using data from ethnographic study of model-building practices in two systems biology labs. Matteo Colombo (Tilburg University, Netherlands) discussed Bayesian cognitive science, inference to the best explanatory framework, and the value of specialization. Peter Sozou (University of Liverpool, UK) argued that a useful theory in cognitive science need not to be a faithful representation of a cognitive process in all its details—instead it should represent some aspect of human cognition in a way that coherently explains a set of observations and ideally leads to testable predictions.

Paul Thagard (University of Waterloo, Canada), the second keynote speaker of the workshop, argued that scientific creativity results from three fundamental computational mechanisms: neural representation, recursive binding and interactive competition. Piotr Giza (Marie Curie-Sklodowska University, Poland) explores possible influences that recent developments in the field of a branch of AI called Automated Discovery Systems may have upon some aspect of the old debate between Francis Bacon's inductivism and Karl Popper's falsificationism. Next, Witold Hensel (University of Bialystok, Poland) discussed Paul Thagard's argument for the approximate truth of deepened scientific theories. Krystyna Bielecka (University of Warsaw, Poland) introduced and discussed Ludwik Fleck's account of science to reconstruct a framework for cognitive research on science.

Hugo Mercier (University of Neuchatel, Switzerland), the next keynote speaker, argued that scientific reasoning does not differ substantially from everyday reasoning, and that the argumentative theory of reasoning can provide a satisfactory account of scientific reasoning. Renne Pesonen (University of Tampere, Finland) discussed the model of intuitive faculty that is best seen as an adapting mechanism that tracks pragmatically relevant features in specific problem instances. Maarten Boudry (Ghent University, Belgium) modelled the cultural epistemology of bona fide science and pseudo-science, drawing on cognitive research on the roots of irrational beliefs and the institutional arrangements of science. Next, Konrad Talmond-Kaminski (University of Finance & Management, Poland) explored the role that epistemic vigilance plays both in scientific and religious cognition. Lucas Afeltowicz (Nicolaus Copernicus University, Poland) discussed the cognitive advantages of two different styles of action: "aristocratic" and "craftsmanship". Finally, Adam Toon (University of Exeter) examined (within the con-

text of situated cognition) Ian Hacking's arguments on different styles of reasoning within scientific practice.

Jakub Ryszard Matyja
Marcin Milkowski
Philosophy, Polish Academy of Sciences

Hypothetical Reasoning, 23–24 August

A conference on "Hypothetical Reasoning" took place at the University of Tübingen in conjunction with ESSLLI 2014. It focussed on the logical aspects of hypothetical reasoning or reasoning under assumptions, which is a key concept of logic, philosophy of science and mathematics. It was organised as part of the French-German ANR-DFG project HYPOTHESES.

The conference started with a keynote talk by Francesca Poggiolesi on counterfactual logics. She presented natural deduction and sequent calculi which are sound and complete for Nute semantics and discussed proof-theoretic results about them. Sergey Melikhov presented a logic of problems and propositions. Torben Braüner discussed Seligman-style deduction for hybrid modal logic and explained how hypothetical reasoning taking place at a particular time can be handled in this setting. Grigory Olkhovikov spoke about truth-value gaps and paradoxes of material implication; he discussed a suitable three-valued logic for which he gave a complete and consistent Hilbert-style axiomatisation.

Paul Egré gave a keynote presentation on negating indicative conditionals. He discussed three kinds of negation and presented results of experimental studies concerning the denial of indicative conditional sentences in natural language dialogues, intended to show how to derive all three negations from a common semantic core, together with additional pragmatic assumptions. Michael Cohen then argued for a new way to understand the notion of 'extending explanatory success' for cases of negative explanatory power.

The first day concluded with two keynote talks. Zoran Petrić discussed representations of multiple-conclusion cuts in proofs by graphs, defined inductively and also purely combinatorially. Finally, Kosta Došen addressed the question of what deductions are. He argued that it is not enough that premisses and conclusions in deductions are propositions, but that it is important that deductions make structures, which in mathematics can be found in categories.

The morning session of the second day was dedicated to hypothetical reasoning in philosophy of science and mathematics. Michel Bourdeau gave a keynote talk on Comte's *Théorie fondamentale des hypothèses*, where he pointed out the historical as well as the conceptual importance of Comte's theory with respect to the question of what kind of hypothesis is admissible. Erdinç Sayan proposed an analysis of the meaning and the confirmation conditions for laws which contain idealisations. Guillaume Schlaepfer discussed hypothetical reasoning and explanations in scientific modeling. Reinhard Kahle explained how the understanding of the notion of axiom has changed in Hilbert, and he illustrated how the reading of axioms as hypotheses can give rise to a

proof-theoretic semantics of non-logical axioms.

Andrzej Indrzejczak gave a keynote talk on hypersequent calculi and linear time; he presented new hypersequent proof systems and results for monomodal and bimodal (temporal) logics of linear frames. Nissim Francez discussed a revision of the notion of harmony in the context of proof-theoretic semantics. Paolo Pistone presented new results on second-order logic.

The conference concluded with a keynote presentation by Arnon Avron. He discussed a consequence relation for sequents, where a sequent is inferred from a set of assumptions which are again sequents, and gave examples for the usefulness of such a relation. A strong cut-elimination theorem for derivations of sequents from assumed sequents was given and its proof explained.

Throughout the conference there were lively discussions. The conference was well attended, including several young researchers and students. Abstracts can be found on the conference website. The proceedings will be published as an open access publication by the end of this year.

Thomas Piecha
Department of Computer Science, Tübingen

Calls for Papers

Euler and Venn Diagrams: special issue of the *Journal of Logic, Language and Information*, deadline 15 October.
Maximum Entropy Applied to Inductive Logic and Reasoning: special issue of *Entropy*, deadline 1 December 2014.
Combining Probability and Logic: special issue of *Journal of Applied Logic*, deadline 15 January 2015.
Causation and Mental Causation: special issue of *Humana.Mente*, deadline 15 March 2015.

What's Hot in ...

Uncertain Reasoning

Andrei Nikolaevich Kolmogorov (1903–1987) is widely recognised as one of the great universal mathematicians of the twentieth century. A glimpse at the table of contents of E. Charpentier, A. Lesne and N. Nikolski (2004: *Kolmogorov's Heritage in Mathematics*, Springer) will certainly persuade potential sceptics that this is far from being an overstatement. To the readers of *The Reasoner*, the name of Kolmogorov is perhaps most readily associated with his axiomatisation of the concept of probability. Indeed the "Kolmogorov axioms" put an end to a controversy over the mathematical status of probability which had lasted for at least two centuries, thereby paving the way for the development of the mathematical theory of probability. G. Shafer and

V. Vovk (2006 "The Sources of Kolmogorov's *Grundbegriffe*", *Statistical Science*, 21(1), 70–98) provide an erudite account of the context and the impact of Kolmogorov's contribution.

Perhaps less well-known to the wider audience is Kolmogorov's contribution to the development of intuitionistic logic. In a seminal paper titled *On the principle of the excluded middle*, which appeared in Russian in 1925 (and which was translated into English only in 1967) Kolmogorov puts forward an axiomatisation of what is now known as *minimal* logic—a proper subset of propositional intuitionistic logic—which anticipated Heyting's 1930 definitive axiomatisation. Seven years later, Kolmogorov wrote a paper in German *On the interpretation of intuitionistic logic* in which he put forward the "proof interpretation" of the connectives—the central feature of what is now known as the Brouwer-Heyting-Kolmogorov (BHK) interpretation. Interested readers are referred to van Dalen, D. (2004: "Kolmogorov and Brouwer on constructive implication and the Ex Falso rule", *Russian Math Surveys*, 59, 247–257) for a comprehensive list of references.

Given his fundamental contribution to the logic of constructive reasoning, it is quite interesting to speculate about what would Kolmogorov think of the fact that the obvious translation of his axioms in a logical language leads to satisfying the *probabilistic excluded middle*, i.e., the fact that for any sentence θ of the classical propositional calculus, the probability of the sentence/event $\theta \vee \neg\theta$ must equal 1. On the other hand, if we give up the law of the excluded middle, fields of sets are no longer complete Boolean algebras. So, at least conceptually, it seems as if Kolmogorov was happy with the algebraic version of the principle and unhappy with its natural logical counterpart.

One rather immediate way out of this alleged *impasse* is to note that Kolmogorov had a frequentist taste for probability, so it seems plausible to think that, to him, the probabilistic version of the excluded middle had little, if anything at all, to do with rational reasoning. Still, it is indeed quite surprising that Kolmogorov appeared not to have any interest in investigating a constructive notion of probability.

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

Introducing . . .

EBM+

Causal claims are crucial in medicine. Bugs, injuries and environmental factors cause disease and other symptoms; medicines, other treatments and public health policies alleviate or prevent such problems. Evidence-based medicine (EBM) is a collection of methods for evaluating the evidence for and against causal claims like these. It provides grading systems and hierarchies of evidence, to help weigh up the evidence and to help decide whether there is sufficient evidence to establish a causal claim.

EBM is good at weighing statistical evidence of associations. Statistical trials are used to test whether there is an association between the putative cause and effect. These

trials vary in size and methodology, and EBM has developed ways of ranking these statistical studies.

Evidence of mechanisms also plays an important role in establishing causal claims. Such evidence is often crucial when it comes to devising a statistical study; interpreting its results; deciding whether an association is causal, due to some other sort of connection, or a statistical blip; or applying the results of a study to a new population or a particular individual. High quality evidence of mechanisms can be produced by statistical studies, but it can also come from literature searches, one-off experiments, imaging, simulations etc. At the moment, EBM does not take explicit account of the role of non-statistical evidence of mechanisms. Arguably, though, we need to take all relevant evidence into account, not just statistical evidence of associations.

EBM+ is a consortium whose members are keen to develop the methods of evidence-based medicine to handle evidence of mechanisms in addition to evidence of associations. The list of current consortium members can be found here. The consortium is keen to welcome new members working in health care or health care methodology who are interested in furthering the aims of EBM+. To sign up, or for more information, please e-mail Michael Wilde, the consortium coordinator.

In the meantime, you can check out the EBM+ blog. Some recent posts include *What's the difference between data and evidence* by Brendan Clarke, *The difference between an ontological and an epistemological question* by Federica Russo, and a post from Andy Fugard on *How we can be fooled into thinking a psychological therapy is effective when it's not*. You can also follow EBM+ on twitter.

Michael Wilde
Philosophy, Kent

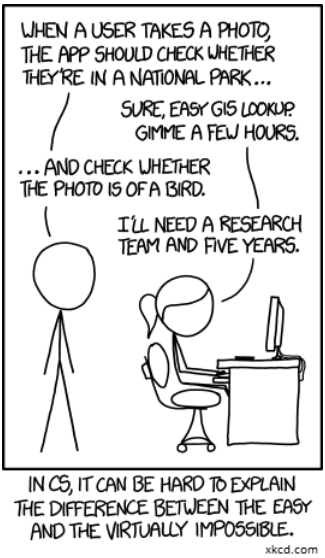

		9	5		6			
		6			7		9	2
	2	7						5
				8	9	1		6
			2		1			
7		3	6	5				
2						9	5	
3	6		8			2		
			4		2	3		

Events

October

WCPA: Western Canadian Philosophical Association, Vancouver, BC, 3–5 October.
FPMW: 6th French Philosophy of Mathematics Workshop, Toulouse, 9–11 October.
Descartes Lecture: Leitgeb on Rational Belief, Tilburg University, Netherlands, 20–22 October.
EBC: Explanantion Beyond Causation, LMU Munich, 23–24 October.
ISR: Inconsistency in Scientific Reasoning, Ghent, 24 October.
ILCS: Inductive Logic and Confirmation in Science, University of Utah, 24–25 October.
Rotman: Knowledge and Models in Climate Science, London, Ontario, Canada, 24–26 October.
ICSR: Knowledge Representation and Reasoning in Robotics, Sydney, Australia, 27–29 October.
MDAI: Modeling Decisions for Artificial Intelligence, Tokyo, Japan, 29–31 October.
IDA: 13th International Symposium on Intelligent Data Analysis, Leuven, Belgium, 30 October–1 November.

November

ECSI: European Conference on Social Intelligence, Barcelona, Spain, 3–5 November.
PoCE: Phenomenology of Cognitive Experiences, University College Dublin, 5–7 November.
Epistemic Reasons: University of Sherbrooke, Canada, 7–8 November.
Grounded Cognition: Düsseldorf, 7–8 November.
ACGC: 8th Arché Graduate Conference, University of St Andrews, 8–9 November.
BotB: Bayes on the Beach, Queensland, Australia, 10–12 November.

LORENTZ: Logics for Social Behaviour, Leiden, 10–14 November.
SoPhiSci: Social Philosophy of Science, Moscow, Russia, 18–19 November.
Epistemic Consequentialism: London School of Economics, 21 November.
ARE& W: Analogical Reasoning East and West, Heidelberg, 24–25 November.
AIC: 2nd International Workshop on Artificial Intelligence and Cognition, Turin, Italy, 26–27 November.
Skepticism: Bonn, 26–28 November.
AAL: Australasian Association for Logic Annual Meeting, University of Canterbury, Christchurch, New Zealand, 29–30 November.

December

NZAP: University of Canterbury, New Zealand, 1–5 December.
Frege: University of Bergen, Norway, 5–6 December.

FE & RE: Formal Epistemology and Religious Epistemology, Oxford University, 8–9 December.
ASCS: Australasian Society for Cognitive Science, Monash University, 8–10 December.
LPMP: Logic and Philosophy of Mathematical Practices, Brussels, 11–12 December.
ABM: Agent-Based Modeling in Philosophy, LMU Munich, 11–13 December.

January

ICLA: 6th Indian Conference on Logic and Its Applications, Bombay, 5–8 January.
DATA: Workshop on the Theory of Big Data Science, University College London, 7–9 January.
ICAART: 7th International Conference on Agents and Artificial Intelligence, Lisbon, Portugal, 10–12 January.
SoTFoM: Competing Foundations, London, 12–13 January.
What is Expertise?: Münster, Germany, 12–13 January.
SAPS: 4th South African Philosophy of Science Colloquium, Pretoria, 15–16 January.
CGCPML: 8h Annual Cambridge Graduate Conference on the Philosophy of Mathematics and Logic, St John's College, Cambridge, 17–18 January.
Diagrams: 1st Indian Winter School on Diagrams, Jadavpur University, Kolkata, 27–31 January.
SDSS: Scientific Discovery in the Social Sciences, London School of Economics, 30–31 January.

Courses and Programmes

Courses

AAAI: Texas, USA, 25–29 January.
Combining Probability and Logic: University of Kent, 20–21 April.
EPICENTER: Spring Course in Epistemic Game Theory, Maastricht University, 8–19 June.
EPICENTER: Mini-course on Games with Unawareness, Maastricht University, 22–23 June.

Programmes

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

Master Programme: in Artificial Intelligence, Radboud University Nijmegen, the Netherlands.
Master Programme: Philosophy and Economics, Institute of Philosophy, University of Bayreuth.
MA in Cognitive Science: School of Politics, International Studies and Philosophy, Queen's University Belfast.
MA in Logic and the Philosophy of Mathematics: Department of Philosophy, University of Bristol.
MA Programmes: in Philosophy of Science, University of Leeds.
MA in Logic and Philosophy of Science: Faculty of Philosophy, Philosophy of Science and Study of Religion, LMU Munich.
MA in Logic and Theory of Science: Department of Logic of the Eotvos Lorand University, Budapest, Hungary.
MA in Metaphysics, Language, and Mind: Department of Philosophy, University of Liverpool.
MA in Mind, Brain and Learning: Westminster Institute of Education, Oxford Brookes University.
MA in Philosophy: by research, Tilburg University.
MA in Philosophy, Science and Society: TiLPS, 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 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.

MSc in Cognitive & Decision Sciences: Psychology, University College London.
MSc in Cognitive Systems: Language, Learning, and Reasoning, University of Potsdam.
MSc in Cognitive Science: University of Osnabrück, Germany.
MSc in Cognitive Psychology/Neuropsychology: School of Psychology, University of Kent.
MSc in Logic: Institute for Logic, Language and Computation, University of Amsterdam.

MSc in Mind, Language & Embodied Cognition: School of Philosophy, Psychology and Language Sciences, University of Edinburgh.
MSc in Philosophy of Science, Technology and Society: University of Twente, The Netherlands.
MRes in Cognitive Science and Humanities: Language, Communication and Organization: Institute for Logic, Cognition, Language, and Information, University of the Basque Country (Donostia San Sebastián).
Open Mind: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.

Jobs and Studentships

Jobs

Post-doc position: in Philosophy, UNAM, deadline 3 October.
Thomas Bayes Chair: of Statistics, Edinburgh, deadline 13 October.
Assistant Professor: in Logic, Algorithms, or Graph Theory, Technical University of Denmark, deadline 15 October.
Assistant Professor: in Philosophy, Department of Philosophy, Logic and Scientific Method, LSE, deadline 31 October.
Assistant Professor: in Philosophy of Science, University of Chicago, deadline 31 October.
Assistant Professor: in Philosophy of Mind, University of Toronto, deadline 13 November.

Studentships

PhD Position: in epistemology and philosophy of science, University of Kent, until filled.
PhD Positions: in "Scientific Realism and the Quantum", Philosophy, Leeds, until filled.
PhD Position: in Spatial Cognition and Reasoning, Psychology, Giessen, until filled.
PhD Position: in Logic and Formal Argumentation, Ruhr-University Bochum, deadline 15 October.
PhD position: in Philosophy of Science, University of Bristol, deadline 20 October.
PhD Position: in Logic and Verification for AI, Utrecht University, deadline 20 October.
PhD Position: in Computational, Mathematical or Philosophical Logic, University of Pretoria, deadline 30 October.
PhD Position: on the project "Changing your mind by changing your brain," Philosophy, Macquarie University, deadline 31 October.
PhD Position: on the project "Recognizing Trust in Natural Language," Computer Science, Philosophy and Linguistics, University of Dundee, deadline 30 November.