
THE REASONER

VOLUME 7, NUMBER 2

FEBRUARY 2013

www.thereasoner.org

ISSN 1757-0522

CONTENTS

Editorial

Features

News

What's Hot in ...

Events

Courses and Programmes

Jobs and Studentships

- 15 sometimes (!) getting better perspective on the more fundamental issues. For example, one session might be devoted to the question of representing the output of ensembles of complex climate models, while the next is devoted to the logic of rational choice when the uncertainty about outcomes is represented as sets of probabilities rather than the standard precise probability function... It is very useful for all participants to engage with the various facets of the broader climate policy problem.
- 15
- 20
- 21
- 22
- 25
- 26



EDITORIAL

It is a pleasure to have another opportunity to guest edit *The Reasoner*, and on this occasion to introduce [Dave Stainforth](#), who is a climate scientist and officially a Senior Research Fellow at the [Grantham Research Institute on Climate Change and the Environment](#) at LSE.

Dave was a prime target as interviewee because he works on a topic that matters a lot to me and to others—the study of climate change—and he clearly brings to this work a great deal of wisdom and genuine concern for doing good and useful science. Dave also has a wonderful infectious energy that is inspiring to others, not least his colleagues.

One initiative of Dave's that I have personally been involved in is the monthly meeting of the 'Climate Change Decision Theory Group' at LSE, which involves climate scientists and statisticians, economists, philosophers, and policy experts. The group discusses the recent work / new ideas of members, where the atmosphere is one of trying to communicate across disciplinary boundaries and, as a result, at least

Dave is a great leader for the CCDTG group because he is interdisciplinary in an effortless way—driven simply by curiosity and concern about problems that cross disciplinary boundaries. This interdisciplinary engagement is also evidenced in Dave's various projects and collaborations. Indeed, I look forward to engaging with more of the work and activities initiated by Dave's group in the Grantham Research Institute.

KATIE STEELE

Philosophy, Logic and Scientific Method, LSE

FEATURES

Interview with Dave Stainforth

Dave and I hang out one Wednesday afternoon in a coffee shop near the LSE to discuss the working life of a climate scientist, with its peculiar methodological concerns and special challenges of faithfully representing knowledge and uncertainty. We start off, however, with the basics...

Katie Steele: From my vantage point, it seems that practising climate scientists require a whole bundle of skills—specialist

knowledge in physics, data analysis and statistics, computer programming, knowledge of policy... Would you say that the field is interdisciplinary in this sense?

Dave Stainforth: Let me answer this question by referring to my own academic background. I came into climate science as a physicist, after university... and I also liked computer programming on the side, sure. (I still enjoy playing with computers and do some programming here and there—it is fun problem-solving and doesn't require too many difficult decisions.) But I started working in climate science very much as a physicist—an atmospheric physicist. In the last ten years, however, I have broadened my interests to communication. I am now very much concerned about the communication of findings that are relevant for assessing impacts of climate change, questions of adaptation and so on. This does benefit from knowledge in a range of areas: physics, numerical analysis, economics and philosophy, to name a few.



I now think it is important to recognise that climate science is a matter of societal concern and is highly policy-relevant. One can choose to simply try to understand climate systems as best as possible, or one can choose to also focus on the social impacts of climate science.

KS: What is the significance of your emphasis on policy here? Are you saying that you now devote a lot of time to science communication, or, rather, that you approach your work as a climate scientist in a different way, i.e., with an eye to policy relevance?

DS: The latter. The attention to policy has led to a shift in emphasis in my scientific work—from modelling and running simulations to the proper interpretation of the data output of these model simulations.

In the past I set up and did a lot of runs (simulations) of these large complex climate models called General Circulation Models [GCMs]. This involves a lot of time and a lot of hard work in getting these computer models up and running... these simulations are difficult to produce. But I have done my time in this respect. The climateprediction.net project that I was involved in is still running, and that's great, but it is up to others now to facilitate the simulations.

The important issue for me now is this: these climate model simulations produce vast output, and there are so many questions about how to analyse these big data sets... In short, what does it all mean? Why run these simulations? We need to really think about what we can get out of these climate models and how the results should be presented. It is tempting to just keep making the models more and more complicated and apparently derive more and more detailed predictions of the type that policy-makers want. Moreover, the power of computers has its own allure... such shiny sophisticated machines that seem to offer endless opportunities for fast and powerful problem-solving... for the mathematically-minded, there is a temptation to create more and more complicated models. We need to be very careful, however, about faithfully representing what we actually *know* about the future climate on the basis of model simulations.

Climate modelling has undergone many advances, but it is my view that we now need to pull back a bit on the modelling exercise and return to fundamental science so as to better understand the use of these climate models and what they can tell us about reality.

KS: I see. So what do you think the models *can* be used for? Do they yield predictions? Is the 'pulling back' just a matter of being more modest about the *precision* of the predictions that we can derive from climate models? Or should we use climate models in quite a different way altogether?

DS: I think it is best to think of climate models as *research tools*—they are useful for understanding interactions between different parts of the climate system. Of course, we don't want to give up on predicting the climate, but we need to be realistic about climate prediction, and in particular, multi-decadal climate prediction, which is of interest to policy-makers. Climate models should be seen as just one of the inputs that allow us to formulate scenarios of how the climate could change in response to different *forcings*. [Forcings are external forces that change the dynamics of the system; a prominent forcing is increased carbon dioxide emissions.] We should aim to formulate scenarios that collectively tell us how the climate could change over time, and give us a general indication of the sensitivity of the climate system.

The main point here is that the output of climate models should not be taken at face value—as predictions of future climate—and presented in more or less unadulterated form to the public and to policy-makers.

KS: Has this been common practice—reporting the direct output of climate models as predictions of future climate?

DS: The output of climate models is generally summarised in colour-coded regional maps indicating changes in climate variables. These are often taken as predictions even when they are not presented as such. This is very common.

But we need to think more about interpreting the data. The vast datasets produced by climate models require interpretation (beyond mere summarising) by climate scientists who can bring to bear background knowledge of climate physics.

KS: Let us pursue this question of understanding the uses and limits of climate models. An interesting question is the relationship between simple and complex models. Might it be the case that for certain coarse variables like global mean temperature, the simpler Energy Balance Models [EBMs] are more telling than the very fine-grained and complex General Circulation Models [GCMs]? Why might that be?

DS: That is a tricky question to answer. But let me say a few things: To begin with, simple models make it easier to understand certain dependencies amongst climate variables. For instance, take global mean temperature. Over the next 100 years, this will mainly depend on forcings, feedbacks within the system and thermal inertia. Now complex systems include all these features, or may include them, but not necessarily as single parameters that may be altered directly by the modeller. Rather, some of these features *emerge* in the model simulations. It is then difficult to determine the robustness or reliability of the final global mean temperature result with respect to these emergent features. In other words, for the complex GCMs, it is very difficult to check the sensitivity of the global mean temperature result with respect to changes in the important variables. The simpler Energy Balance models, by contrast, do allow the modeller to adjust these variables.

GCMs, on the other hand, are useful for other things. They

are by no means faithful representations of the actual climate system, but they do aim to represent a physical climate and can help us better understand how feedbacks interact; GCMs inform us about interactions in the climate system and this helps us to understand the simpler models and what are realistic relationships between the main variables in the simpler models.

KS: So you are suggesting that the complex General Circulation models may be useful for investigating major climate processes even if they do not necessarily yield accurate predictions for variables like temperature in some region at some snapshot in time?

DS: It may be futile to try to put too much realism into a climate model in an effort to increase predictive power. But different models may be useful for investigating different processes. For instance, we may have one model that includes modules like a complex land-surface scheme, such that we could, in principle, detect a phenomenon such as the dieback of the Amazon rainforest. This is not necessarily a realistic example, but from this model, say, we may be able to learn how ocean circulation patterns affect / are affected by Amazon dieback. But then if we are interested in the question of how changes in the stratosphere are related to changes in the troposphere, we may not need the complex land-surface scheme in our model. A different model might be more appropriate for examining the stratosphere-troposphere interaction.

There is a tendency to try to represent more and more aspects of climate in one model. In the light of model error, it is not clear that this necessarily leads to better predictions, and given the danger of over-interpreting the results of climate models, the aspiration to simulate everything together seems dangerous. Better to try to come to terms with the fact that different models are useful for understanding different climate processes, I would say.

KS: And what about verifying these models? How can we verify and improve the models if they are more about understanding processes than yielding realistic predictions? When it comes to model improvement, is it then just a matter of continuing to search for insights from basic physics?

DS: Well, yes, keeping on thinking about the physics is a good idea.

Now, there is a tendency to use historical records of, say, temperature, for confirmation. This shows that the community aspires to models that directly predict such variables at a fine-grained level, and does not just aspire to understanding processes. As we have discussed, I am not sure that this sort of predictive power is a realistic possibility. In any case, the data is constituted by the actual temperature record of the last 50 years or so, and further temperature data from the past that is inferred from various proxy data (e.g., ice cores). We can work out what models best fit the past data, in line with the way we go about constructing and improving weather models. We may even try to reconstruct / explain past climatic events, i.e., hypothesise what were the circulation patterns that lead to those events. This is all very well. The problem is that these fitted models of long-term climate cannot be periodically tested in the way that short-term weather models can. Moreover, future climate dynamics are expected to be quite different from past climate dynamics due to the various forcings on the system.

So let us return to the more modest notion that models may increase our understanding of climate processes. Perhaps another way to verify the models on this score, besides consulting basic physics, is to check them against each other. We can

check whether, for instance, the relevant models agree on the broad causes and effects of, say, Amazon dieback. . . Of course, this raises the question of what is the significance of models agreeing on some finding, and how this should affect our confidence in this finding. If the models agree on the process connections that lead to certain events then maybe this can point the science in a useful direction: to physically understand those connections would represent an improvement in scientific understanding and provide some form of confirmation of this aspect of the models.

KS: Yes, the question of the epistemic significance of model agreement is tricky. In general, a bunch of evidence provides greater confirmation of a hypothesis (such as, for instance, the causes of Amazon dieback) the more independent and the more diverse the evidence, all other things being equal (chiefly, the reliability of each piece of evidence in isolation).

DS: It is a very difficult question as to what extent climate models can be classified as independent. Many share large sections of computer code. And even if they have been coded differently, they will to some extent share dynamic equations and assumptions from physics. . . The question is the extent to which this compromises independence. . .

KS: Maybe it helps to think in terms of witnesses to a crime. Of course, if all the witnesses say the person at the scene of the crime was wearing a blue shirt, it would not mean the witnesses are not independent. They only need be independent conditional on the truth of the matter. In the witness case, we tend to think they are independent if they arrive at their conclusions using their own individual sensory apparatus. . . Well, perhaps what is important is that their errors of judgment are independent. If we have reason to think that the witnesses all make the same kinds of vision mistakes in the same kinds of circumstances, then their reports would not be independent.

So could we expect the errors of climate models to be independent? I guess to answer this question we come back to consulting basic physics.

DS: Yes it is a tricky issue. One thing that deserves more attention is the study of climate within a model itself. My worry is that oftentimes we don't even understand the model climate systems, let alone how the model climate relates to the real world. We need to thoroughly test processes within the climate model—what possible circumstances can bring about an Amazon dieback, for instance, and does the dieback always have the same effects?

KS: Actually, speaking of 'climate within a model' . . . What *is* climate? I forgot to ask this question!

DS: Surprising as it may be, this is not a trivial question and you will have to stay tuned for a forthcoming paper that I am writing on this topic with my colleague [Lenny Smith](#). Climate is typically understood as some sort of statistical summary of weather variables. This in itself leaves many questions open about what sort of summary data we are interested in and what sort of timeframes for these statistical summaries are appropriate. There may be more scientifically fruitful ways, however, to understand the term 'climate'. I will leave this question hanging for now.

KS: Sounds interesting. Let us now finally turn to a question that is not so much about climate science itself, but rather about communicating climate science to the greater public and to policy-makers. I have heard you express the view that the focus on 'scientific consensus' in the public debate is unhelpful, and perhaps even damaging to public understanding of climate

science. Why do you hold that view?

DS: I think it is unhelpful because people realise that there is a contradiction in talking about a consensus—it seems to imply that everything is understood and no-one disagrees with each other, when just a little amateur scientific research reveals that this is not so. Of course, climate scientists are not disagreeing about the basic fact that climate change is a very major concern, and that increased greenhouse gases leads to warming. But they disagree on so much else that talk of ‘consensus’ seems a falsehood.

Moreover, the notion of ‘consensus’ suggests something prescriptive—that there should be no arguments. But in fact: there should be more public arguments amongst scientists. In my opinion, the more apparent the arguments amongst scientists regarding the wheres and whys of the details of climate science, the more weighty the issue will appear to be. The well understood fact that climate change is a big concern will shine through in these arguments.

KS: I think this is a nice point. But this focus on consensus was itself a response to public distrust of climate science. Why do you think there has been such distrust of climate science in particular? Do you think it is simply because the science is very close to policy, and suggests quite big policy changes?

DS: Yes, I do think that is the reason for the difference between the response to climate science versus other sciences.

KS: In any case, your view that it is better to communicate the uncertainties and methodological issues associated with climate science was affirmed in my opinion by the reaction to the exhibit you organised at the [Royal Society Science Fair](#) last year. Initially I thought you were being rather ambitious in trying to communicate issues of risk versus model uncertainty, the status of predictions from ‘model ensembles’, and problems of decision-making under severe uncertainty, but then I saw that people really engaged a lot more with the climate change predicament when they were brought into the ‘inner circle’ with respect to these issues.

DS: Indeed. That was the idea. And of course, that highlights another reason not to focus on consensus in climate science—it hides all the interesting and difficult problems. We are talking about a relatively new area of science, and of course an important area of science, and we want to spark intellectual curiosity on these topics, and to enthuse more people to work on these problems.

KS: Well this call to arms is I think a good note to end on... Thanks very much Dave for taking the time out for this chat today!

Is Ethical Relativism Self-Stultifying?

Ethical relativism is purported to maintain there is no single true morality and thereby to encourage an attitude of tolerance or non-interference *vis-à-vis* competing and incompatible bodies of moral value. (See, David Wong 1984: *Moral Relativity*, University of California Press.) However, critics of ethical relativism have found the combination of these two propositions incoherent and self-refuting. Bernard Williams, for example, argues that ethical relativism consists of three propositions: (1) ‘right’ means ‘right for a given society’, (2) ‘right for a given society’ is understood in a functionalist sense, and (3) it is wrong for people in one society to condemn or interfere with the values of another society. But, ‘the view is clearly inconsistent since it makes a claim in its third proposi-

tion, about what is right and wrong in one’s dealings with other societies, which uses a *nonrelative* sense of ‘right’ not allowed for in the first proposition’ (Bernard Williams 1980: *Morality*, Cambridge University Press, p. 34). Thus ethical relativism is a ‘logically unhappy attachment of a nonrelative morality of toleration or non-interference to a view of morality as relative’ Williams (1980: p. 35).

But, is ethical relativism guilty of logical inconsistency? The purpose here is to propose a way of exonerating ethical relativism from the charge of incoherence by drawing on Alfred Tarski’s distinction between *object-language* and *meta-language* in ‘The Semantic Conception of Truth’ (Alfred Tarski 1944: ‘The Semantic Conception of Truth,’ *Philosophy and Phenomenological Research* 4, 341–376). In discussing the problem of defining truth against the backdrop of paradoxes like the liar antinomy, Tarski suggests that ‘we have to use two different languages’: the first is the language which is “talked about” and the second is the language in which we “talk about” the first language (Tarski 1944: p. 349). In this dichotomy, statements involving the concept of truth are strictly speaking not uttered at the same level of language use and should be lasoed into two different levels of object and meta utterances.

Applying this Tarskian schema to the pronouncements of ethical relativism, the claim is basically that the doctrine of ethical relativism (1) and the doctrine of tolerance (3), are not at the same level of language use and thus their utterance should be bifurcated into two different levels of linguistic expression. Should this separation of linguistic levels of utterance work, it would show that ethical relativism is not self-stultifying after all. The application of the Tarskian schema to ethical relativism may run thus: having observed the variety, variability, incompatibility, and incommensurability of ethical statements made by people across different cultures and different times or that they were just hypothesized and imagined, the ethical relativist reaches her position that there is no single true morality. On this model, ordinary moral judgments are at the *object* level as they are on a par with all the other moral judgments made by individuals in other societies and cultures across time and space. Thus, common-or-garden moral statements are uttered at the same level of language that is “talked about”. Yet, the ethical relativist’s own pronouncement is at a *meta* level of language use since she is “talking about” the language in which ordinary moral agents “talk about” their moral judgments.

So far there is no logical inconsistency, but neither is there any talk of tolerance. Thus, where does tolerance fit into this model without degenerating into self-contradiction? Is not tolerance at the *same* meta-level of language use as the doctrine of ethical relativism itself, thereby threatening the cogency of the position? Indeed, this is exactly the claim that Williams makes when he puts the ethical relativist’s first and third claims on a par. But, are they at the *same* level?

At this juncture, the ethical relativist seems to have two options to shun the snare of logical inconsistency. (A) She may concede that claims (1) and (3) are at the same meta-level but deny that they belong to the same category of statements. That is, not all meta-level statements are about the same subject: although both claims are normative by nature, the first one has its normativity rooted in ethics whereas the second one is anchored in epistemology. Tolerance is an *epistemic* norm given that from an epistemological perspective there is no way to privilege one moral decision making over another. The norms involved in the two statements belong to different species of

normativity, and thus the *relativity* of one is not in conflict with the *non-relativity* of the other.

However, if the distinction between different types of normativity fails to have purchase on hardcore non-relativists, the ethical relativist may resort to a second option (B): unlike option (A), it is denied that propositions (1) and (3) are at the same level of language use. The ethical relativist reaches proposition (1) at the *meta*-level while observing how individuals in various cultures express their moral judgments at the *object*-level. However, we are still none-the-wiser about the moral relativist's *normative* ethics: that is, what is her moral assessment of those variable, incompatible, or incommensurable ethical values? Given her meta-ethical position, normatively she can neither condone nor condemn any particular moral value. Her meta-ethical stance bars her from celebrating or censuring any value unless she descends to the object-level where she can approve or disapprove of moral values in accordance with the moral values of her culture. But, once she arrives at her meta-ethical position of relativism, she reaches a stalemate in terms of normative assessment of moral values belonging to her or other cultures. It is here that the moral relativist is forced to ascend to the next level of discourse, *viz.*, the meta-meta-level, where she can propose the moral norm of tolerance against the twofold background of her meta-ethical position and the stalemate of not being able to offer a normative judgment on any particular moral value. Should the application of such a Tarskian hierarchy to the moral relativist's endeavor work, then, *contra* Williams, the ethical relativist can consistently maintain her moral relativism with her recommendation of tolerance since the two theses are arrived at two different levels of discourse.

MAJID AMINI

Philosophy, Virginia State

The Contextual Theory of Scientific Understanding: A Rejoinder to Erik Weber

In two papers published in *The Reasoner* (Volume 6, issues 4 and 8), Erik Weber advances several objections against the contextual theory of scientific understanding that was presented in Henk de Regt and Dennis Dieks (2005: "A Contextual Approach to Scientific Understanding", *Synthese* 144: 137–170) and further developed in subsequent publications. Weber claims that our theory confronts two problems and contains two gaps. In this rejoinder I will address Weber's concerns and defend our theory against his criticisms.

The first alleged problem for our theory is that it "is not contextual with respect to the meaning of understanding". As Weber rightly observes, our theory is contextual because we argue that the tools scientists use to achieve understanding vary with the (historical and disciplinary) context and that success in achieving understanding depends on scientists' skills. He objects, however, that we ascribe a context-independent *meaning* to understanding as a general aim of science (p. 61). But why is this a problem? Our theory accounts for important contextual variations in understanding, while retaining the idea of understanding as a universal aim of science. If the meaning of understanding would change with the context, no general, non-trivial account of scientific understanding would be possible. But perhaps Weber objects to what he thinks we propose as the meaning of scientific understanding, namely "qual-

itative derivations". This would be a mistake, however: we do not *define* understanding in terms of qualitative derivation. Criterion CIT, which states that a scientific theory is intelligible if scientists can recognise qualitatively characteristic consequences of it without performing exact calculations, is not intended as a definition but as a *test* for the intelligibility of theories. This idea—which was implicit in our 2005 paper—is explicated in De Regt (2009: "The Epistemic Value of Understanding", *Philosophy of Science* 76: 585–597), where intelligibility is defined as "the value that scientists attribute to the cluster of virtues (of a theory in one or more of its representations) that facilitate the use of the theory for the construction of models" (p. 593). Intelligible theories are required for understanding phenomena, as criterion CUP states: "A phenomenon P is understood *iff* a theory T of P exists that is intelligible (and meets the usual logical, methodological and empirical requirements)" (*ibid.*).

Accordingly, Weber's contraction of our criteria CUP and CIT into QD ("In all possible contexts understanding as an epistemic aim of science consists in the capacity to make qualitative derivations with the theory", p. 61) does not represent our view. CIT is one way of testing the intelligibility of theories, which may fail to apply in some contexts. Thus, understanding a theory is not identified with qualitative derivation. Moreover, QD obscures the relation between the understanding of phenomena and the intelligibility of theories, as specified by CUP.

The second problem Weber sees for our account is that its scope is restricted to cases where theories are used. Citing an example involving the pendulum law, he maintains that in many cases understanding can be achieved without using theories. We think this is moot at best. To be sure, laws and models may provide understanding, but only if they are embedded in, or constructed on the basis of, more general theories. If not, they do not yield understanding. This distinguishes them from cases like the classic barometer example, in which (qualitative) prediction is possible but no understanding gained. Weber's pendulum example is similar: deriving the period from the length and the pendulum law (which only specifies a relation between period and length) does not qualify as understanding, whereas qualitative prediction of the behaviour of a pendulum on the basis of Newtonian theory would qualify as such because it evidences a grasp of how the phenomenon is brought about.

Weber's second paper discusses two alleged 'gaps' in our theory. First, he claims that we do not argue for our thesis that qualitative derivation is crucial to understanding (CIT). "What makes qualitative derivation interesting so that it can count as an additional aim of science?", asks Weber (p. 131). As mentioned above, CIT is not a definition of understanding, and thereby not an aim in itself. It is an indicator of the intelligibility of a theory (for scientists), which in turn is a precondition for the construction of explanations of phenomena (CUP). But this does not imply that it is merely of heuristic value. Explanations (involving quantitative calculations) provide understanding only if scientists have insight into how they are constructed, and this requires intelligibility of the associated theories, which can be tested via CIT. It is in this way that Boltzmann's qualitative account of Boyle's law shows how kinetic explanations actually provide understanding.

The second 'gap', according to Weber, is that it is unclear how our theory of understanding relates to the "complex-system mechanist tradition", defended by Machamer, Darden and Craver, who claim that mechanisms are intelligible because

they describe “bottom-out entities and activities” (p. 131). Our thesis that ways to achieve understanding are context-dependent implies, by contrast, that mechanistic analyses are not intrinsically intelligible, nor essential to understanding. It is a contingent fact that understanding Boyle’s law via the kinetic theory involves a description of mechanisms. Macroscopic gas phenomena can be understood in terms of a microscopic theory but also on their own level, via thermodynamics. To be sure, statistical mechanics explains more phenomena than thermodynamics (e.g., Brownian motion), but this does not entail that it is a more intelligible theory or that descriptions of mechanisms are essential to scientific understanding. There are no essential or in-principle restrictions on the nature of theories used in explanations. Often it is counter-productive to search for microscopic explanations and more fruitful to stay at the macroscopic level. What is more, sometimes a microscopic phenomenon is better understood with a macroscopic theory. For example, one better understands the spreading of smoke particles produced by a cigarette with the second law of thermodynamics than through an exact microscopic description of their individual paths. Accordingly, there is no fundamental *explanatory* asymmetry.

Complex-system mechanists correctly assert that mechanistic accounts often provide understanding. But they are mistaken when they base this on a postulated *intrinsic* intelligibility of mechanisms. Our analysis of understanding, by contrast, provides an argument: it shows *how* mechanistic analyses may (in particular contexts) enhance intelligibility and produce understanding. Thus, our theory is congenial to the mechanistic approach. It rejects the idea of intrinsic intelligibility, but mechanists can well do without.

HENK W. DE REGT

Philosophy, VU University Amsterdam

NEWS

Epistemic Groups and Collaborative Research in Science, 17–19 December

The [conference](#) was held in Nancy (France) in the MSH Lorraine, supported by the Archives Henri Poincaré (CNRS and Université de Lorraine). The topic of the conference was to discuss philosophical issues about epistemic groups, be it in everyday life or in science, and about collaborative research in science, i.e., when several scientists work together to achieve a result. The conference enabled interactions and cross fertilization between approaches or disciplines that tackle this object of inquiry with different methods—philosophy of science, formal epistemology, history of science, social epistemology or psychology. This conference can be seen as a continuation, with a more narrow focus, of a previous conference held in the same place last year (‘The Collective Dimension of Science’, cf. this journal, 6(2) p. 25).

The conference featured six keynote speakers. As a conference opening, Erik Olsson presented a Bayesian model which studies belief updating—and also trust updating—in a social context. This model, called Laputa, can be implemented on a computer, so as to investigate more thoroughly properties of networks, like group polarization of beliefs. K. Brad Wray examined how collaborative research affects the epistemic culture of science. He analyzed for instance the way collective au-

thorship is dealt with in various scientific fields, and discussed specific problems that collaborative works pose to the usual refereeing process. On Tuesday, Bryce Huebner (in a joint work with R. Kukla and E. Winsberg) discussed epistemic accountability in collaborative works, in particular in the light of the various non-epistemic pressures or interests of their members. He suggested that providing a social epistemic model of the collaboration could be part of the solution. Jan Sprenger (based on a joint work with S. Hartmann and R. Dawid) presented a socio-epistemic variant of the No Alternatives Argument—i.e., does the current lack of alternative to a successful research program, call it H, provide a valid argument in favor of H? He also discussed to what extent this argument could be used by the research administration. Finally, Wednesday featured two lectures on judgment aggregation theory. First, Denis Bonnay discussed the attribution of beliefs to non-organized groups, and suggested that the clustering method, currently used in market research, could enrich the classical judgment aggregation theory. Second, Christian List distinguished three kinds of collective attitudes: the aggregate attitude, which is plainly metaphorical and reductionist, the common attitude, in a classical Lewisian flavor, and the corporate attitude, in which the group constitutes an agent in its own right.

Fifteen anonymously selected papers were presented at this conference, which tackled various questions related to epistemic groups and scientific collaboration. They investigated for instance the dynamics of scientific groups, the division of cognitive labor, the decision procedures within groups, the authorship for collaborative works, or the definition of group knowledge. Several talks described cases of scientific collaborations or analyzed the functioning of epistemic groups, in physics, bioethics, environment or computer-modeling. From a methodological viewpoint, the methods used involved game-theory, numerical simulations, multi-agent simulations, formal reasoning, conceptual analysis, or historical and sociological case-studies. Overall, this conference was rich of the diversity of the academic fields of the participants, of the variety of the methodologies used and of the fruitfulness of the discussions.

Proceedings of the conference are planned to appear in a special issue of an international journal or in a book.

THOMAS BOYER

Université Lille 1

Archives Henri Poincaré, Nancy

CYRILLE IMBERT

Archives Henri Poincaré, Nancy

Theoretical Aspects of Rationality and Knowledge, 7–9 January

The 14th Conference on Theoretical Aspects of Rationality and Knowledge ([TARK 2013](#)) was held from January 7 to 9, 2013, at the Institute of Mathematical Sciences, Chennai, India. TARK brings together researchers from a wide variety of fields, including Artificial Intelligence, Cryptography, Distributed Computing, Economics and Game Theory, Linguistics, Philosophy, and Psychology, in order to further the understanding of interdisciplinary issues involving reasoning about rationality and knowledge. For the first time, the conference was held in India and in January.

Similar to previous TARK conferences, out of 64 submissions to the conference, 18 papers were accepted as contributed

talks and 8 as poster presentation by an interdisciplinary program committee chaired by Burkhard C. Schipper, University of California, Davis. The program included work by Adam Bjørndahl, Joe Halpern and Rafael Pass on [language-based games](#). In standard game theory, the domain of the utility functions includes just the outcomes in the game. Language-based games generalize the domain to maximal consistent sets of formulas in some language, thus allowing to model psychological notions like guilt or reference-dependence. It also included a paper by Jayant Ganguli and Aviad Heifetz on [universal interactive preferences](#). In games, a player's belief about other players, beliefs about beliefs etc. are modeled with type spaces a la Harsanyi. Usually, those beliefs are formalized by subjective probabilities or equivalently by subjective expected utility a la Savage. For such beliefs, Mertens and Zamir have previously shown the existence of a universal type space, i.e., a type space that contains all hierarchies of beliefs that appear in some type space. Ganguli and Heifetz generalize this result to preferences that include Choquet expected utility theory, maxmin expected utility, lexicographic expected utility, incomplete preferences etc. Other talks included topics such as coordination and common knowledge, awareness, dynamic epistemic logic, defeasible reasoning, extensions of Aumann's "No-agreeing-to-disagree" result, computational social choice, and more.

The contributed talks were complemented by three invited speakers. Rineke Verbrugge gave an insightful talk about her empirical work on logic especially the "theory-of-mind". Lin Fangzhen made his work on computer-aided theorem discovery in game theory accessible to the TARK audience. Unfortunately, due to illness, Pierpaolo Battigalli had to cancel at the last minute his talk on dynamic epistemic game theory. Burkhard C. Schipper presented instead a survey on epistemic modeling approaches to awareness.

The local organization was extremely well done thanks to the local organizing committee chaired by R. Ramanujam. TARK was co-located with the 5th Indian Conference on Logic and its Applications (ICLA), which was held just right after TARK. Participants were also able to enjoy Carnatic music and Indian dance from the music and dance festivals held concurrently in Chennai.

The proceedings of the conference will appear at [ACM digital library](#). They will become available also from <http://www.tark.org/>.

BURKHARD C. SCHIPPER

Economics, University of California, Davis

Aims and Norms, 18 January

The Aims and Norms: Reasoning workshop took place in the University of Southampton, and was organised by this university in collaboration with the [Centre for the Study of Mind in Nature](#) (Oslo). It was the fourth event in the [Aims and Norms](#) workshop series.

The workshop focused on the norms and aims governing our processes of reasoning. The first speaker was Clayton Littlejohn (King's College London). In his talk, entitled 'Knowledge is Probably the Norm of Belief', he argued against truth as the norm of belief. Littlejohn defended a conformity account of guiding reasons: we are supposed to remain within the guidelines established by normative reasons, but it is not demanded that those are the reasons *for which* we act. If this account were

right, the norm of truth for belief would not manage to explain the inward focus of epistemic appraisal (epistemically, it matters how we arrive at true beliefs)—the norm of knowledge, however, would provide a suitable explanation. Tom Simpson (Cambridge) offered a response to the paper.

The second paper, 'Kant on the Constitutive Aim of Theoretical Reason', was presented by Sasha Mudd (Southampton). She interpreted Kant as arguing that the constitutive aim of theoretical reasoning is the complete systematic unity of our empirical knowledge. Although this is an unattainable target, it may play a regulative guiding role: from this aim, there would follow several norms to the effect of proscribing the decrease of systematic unity in our beliefs. Mudd explored the possibility that this transcendental goal is the aim of our reasoning, and defended it from possible objections. A response to Mudd's paper was provided by Felix Koch (Göttingen).

Kieran Setiya (Pittsburgh) argued in his paper ('Epistemic Agency: Some Doubts') that the notion of epistemic agency only makes sense in a deflationary reading: otherwise, it is mysterious or misguided idea. Setiya claimed that believing is a static standing condition, nor an event or act. Furthermore, he argued that believing that p on the grounds that q , amounts to having the belief that p , and the belief that the fact that q is evidence that p . Thus, inferring is not a dynamic activity in a robust sense. Sophie Edwards (Southampton) commented Setiya's paper.

The final speaker was Anders Nes (CSMN, Oslo); the title of his talk was 'Basing Relations Aim at Normative Reasons'. Nes described the position he defends as reasonism about basing relations: a belief that p is based on another belief that q only when the relation between those beliefs is subject to a norm to the effect that the second belief offers good reasons for the first one. Nes argued that reasonism does not require having a belief about such norm being satisfied—that is, a belief that q provides good reasons for believing that p . The respondent for Nes' paper was Guy Fletcher (Edinburgh).

The workshop was supported by the University of Southampton, the Faculty of Humanities, the CSMN (Oslo), the British Academy, and the Mind Association.

JAVIER GONZÁLEZ DE PRADO SALAS

Departamento de Lógica, Historia y Filosofía de la Ciencia,
UNED

& University of Southampton

Calls for Papers

[THE QUESTION OF BIO-MACHINE HYBRIDS](#): special issue of *Philosophy and Technology*, deadline 28 February.

[HYPERINTENSIONALITY](#): special issue of *Synthese*, deadline 1 March.

[THE SQUARE OF OPPOSITION](#): special issue of *History and Philosophy of Logic*, deadline 30 June.

[INFINITE REGRESS](#): special issue of *Synthese*, deadline 1 July.

WHAT'S HOT IN . . .

Logic and Rational Interaction

It is the time to celebrate a special birthday in this young year: Yablo's paradox is turning 20. In 1993, Stephen Yablo published his seminal paper, '[Paradox without Self-Reference](#)'

which claims to be doing exactly this: provide a liar-type paradox without self-referentiality.

The Yablo set consists of statements Y_i indexed with the natural numbers, where each statement says that all Y_j with $j > i$ are wrong. It is easy to see there is no consistent ascription of truth values to the set of all Y_i .

Yablo himself claims that the paradox does not involve self-referentiality, since none of the Y_i refers, implicitly or explicitly, to itself. However, this question has been the topic of a long and ongoing debate involving, *inter alia*, Yablo, Leitgeb and Priest.

In a recent paper, [Ming Hsiung](#) presents a new connection between Yablo's paradox and the Liar paradox. He shows that both are equiparadoxical, thereby making an argument for Yablo's paradox to rely on circularity.

In another [forthcoming paper](#), Cezary Cieřliński and Rafał Urbaniak examine the behaviour of Yablo's paradox when truth is replaced by provability in a sufficiently strong background theory.

Finally, [Thomas Forster](#) connects a logical analysis of Yablo sequences to the omitting types theorem.

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](#).

DOMINIK KLEIN
TiLPS, Tilburg University

Uncertain Reasoning

Logic is often partitioned into classical and non-classical. Whilst we all have a vague intuition about the meaning of the distinction, it turns out to be quite difficult to pinpoint a precise characterisation of it. The partition is, for a start, hardly symmetrical. Classical logic is always thought of in the singular, whereas there is an intrinsic plurality of non-classical logics. In addition, some non-classical logics—notably intuitionistic logic—take issue with the universal validity of their classical counterpart, whilst other logics are put forward as non-trivial extensions—friendly amendments, that is—of classical logic. As it turns out, some non-classical logics are more classical than others. Modal and many-valued logics are perhaps the names which most readily come to mind in this respect.

One might thus be led to the following consideration. It might well be that practitioners implicitly rank the classicality of logics based on their mathematical depth or relevance. This would account for the fact that non-classical logics are often categorised as “philosophical logic”, as witnessed by the fact that the standard reference for the subject has been, for the past three decades, the monumental *Handbook of Philosophical Logic*.

This way of putting the distinction doesn't really hold water, though. First, non-classical logics featuring in the *Handbook* have long been recognised to be worthy of serious mathematical attention. These have been applied very successfully in a number of areas other than philosophy, including among others computer science and linguistics. A very interesting, albeit perhaps under-cited, position paper on this is Larry Moss (2005: *Applied Logic: A Manifesto* in D. M. Gabbay, S. Goncharov, & M. Zakharyashev, eds, *Mathematical Problems*

from *Applied Logic I*, pp. 317–343, Springer).

Second, non-classical logics and methods may serve as a basis for the development of *non-classical mathematics*. The [February 2013 special issue of the *Logic Journal of IGPL*](#), edited by Libor Behounek, Greg Restall, and Giovanni Sambin offers a rich collection of papers exploring this extremely fascinating idea. The contributions span many areas, from set theory to reverse mathematics to abstract algebra, and covers a number of very interesting foundational questions, most notably that of mathematical pluralism.

The best-known example of non-classical mathematics is constructivism, which developed against the background of intuitionistic logic. Yet, as the editorial to this special issue points out, there are a number of interesting yet less well-known theories which couple classical as well as non-classical logic with non-classical foundations of mathematics:

Just like mathematical theories over non-classical logics, such theories offer a different perspective on certain mathematical objects—in some of them, for instance, the reals are countable—and enable methods alternative to those of traditional mathematics (e.g. non-standard analysis). Finally, also the meta-mathematical study of predicate non-classical logics themselves can be viewed as an indispensable prerequisite to, and so a part of, non-classical mathematics.

HYKEL HOSNI
Scuola Normale Superiore, Pisa
CPNSS, LSE



EVENTS

FEBRUARY

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

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

C&S: Causation and Structuralism Workshop, Cologne, Germany, 8 February.

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

LAFLANG: 2nd International Workshop on Learning, Agents and Formal Languages, Barcelona, Spain, 15–18 February.
ICAART: 5th International Conference on Agents and Artificial Intelligence, Barcelona, Spain, 15–18 February.
CSEE: 2nd International Conference on Advances in Computer Science and Electronics Engineering, New Delhi, India, 23–24 February.
SAPHIR: Systematic Analytic Philosophy and Interdisciplinary Research, Ruhr-Universität Bochum, 25–27 February.
STACS: 30th Symposium on Theoretical Aspects of Computer Science, Kiel, Germany, 27 February–2 March.

MARCH

THEORETICAL AGENCY: Auburn, Alabama, 1–2 March.
PTS: 2nd Conference on Proof-Theoretic Semantics, Tübingen, Germany, 8–10 March.
LKL: Logic, Knowledge, and Language, Paul Gochet Memorial Conference, Brussels, Belgium, 14–15 March.
PHILOSTEM: 5th Midwest Workshop in Philosophy of Science, Technology, Engineering, and Mathematics, Fort Wayne, IN, 14–16 March.
METAPHYSICAL VIRTUES: Western Michigan University, Kalamazoo, Michigan, 15–17 March.
SIMRIDE: 1st workshop on Uncertainty Quantification and Data Assimilation in Numerical Simulation of Physical Systems for Risk-Informed Decision Making, Durham, 18–21 March.
INFORMATION: 5th Workshop on Philosophy of Information, University of Hertfordshire, UK, 27–28 March.
UNILOG: 4th World Congress and School on Universal Logic, Rio de Janeiro, Brazil, 29 March–7 April.

APRIL

SBP: International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction, UCDC Center,

Washington DC, USA, 2–5 April.

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

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

THE ANALYSIS OF THEORETICAL TERMS: Munich, Germany, 3–5 April.

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

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

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

PERCEPTION, MODELS, AND LEARNING: 15th Annual Pitt-CMU Graduate Conference, Carnegie Mellon University, 5–6 April.

ADS: Agent-directed Simulation Symposium, Bahia Resort, San Diego, CA, USA, 7–10 April.

INFORMATION: SPACE, TIME, AND IDENTITY: Milton Keynes, 8–10 April.

PHDs IN LOGIC: Munich, 8–10 April.

MODELS & DECISIONS: 6th Munich-Sydney-Tilburg Conference, Munich, 10–12 April.

IDENTITY AND PARADOX: Lille, France, 11–12 April.

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

IEEE-SSCI: Symposium Series on Computational Intelligence, Singapore, 15–19 April.

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

R&R: Reasons and Reasoning, Georgetown University, 20 April.

IMPLICIT BIAS: University of Sheffield, 20–21 April.

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

GIRL@LUND: 2nd Conference on Games, Interactive Rationality, and Learning, Lund, 23–26 April.

EXPLANATORY POWER: Understanding Through Modeling. Epistemology, Semantics, and Metaphysics of “Inadequate”, Ruhr-Universität Bochum, 25–26 April.

PoM&PSYCH: KCL Graduate Conference in Philosophy of Mind and Psychology, Institute of Philosophy, Senate House, London, 26 April.

PHILOSOPHY OF INFORMATION: The Value of Information, American University, Washington DC, 26 April.

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

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

MAY

ICLR: 1st International Conference on Learning Representations, Scottsdale, Arizona, 2–4 May.

SDM: 13th SIAM International Conference on Data Mining, Austin, Texas, USA, 2–4 May.

O&M: Ontology and Methodology, Virginia Tech, 4–5 May.

CTFoM: Category-Theoretic Foundations of Mathematics, Irvine, California, 4–5 May.

MSDM: 8th Workshop on Multiagent Sequential Decision Making Under Uncertainty, Saint Paul, Minnesota, USA, 6–7 May.

EMAS: 1st International Workshop on Engineering Multi-Agent Systems, Saint Paul, Minnesota, USA, 6–7 May.

ALA: Adaptive and Learning Agents Workshop, Saint Paul, Minnesota, US, 6–7 May.

MSDM: Multiagent Sequential Decision Making Under Uncertainty workshop, Saint Paul, Minnesota, USA, 6–7 May.

AAMAS: 12th International Conference on Autonomous Agents and Multiagent Systems, Saint Paul, Minnesota, USA, 6–10 May.

ADMI: 9th International Workshop on Agents and Data Mining Interaction, Saint Paul, Minnesota, USA, 6–10 May.

PHILANG: 3rd International Conference on Philosophy of Language and Linguistics, University of Lodz, Poland, 9–11 May.

UK-CIM: Causal Inference in Health and Social Sciences, University of Manchester, 14–15 May.

MCS: 11th International Conference on Multiple Classifier Systems, Nanjing University, China, 15–17 May.

MATHEMATISING SCIENCE: University of East Anglia, Norwich, 16–17 May.

LMP: 13th Philosophy of Logic, Math and Physics Graduate Conference, Ontario, Canada, 18–19 May.

SLACRR: St. Louis Annual Conference on Reasons and Rationality, St Louis, MO, 19–21 May.

TAMC: 10th Conference on Theory and Applications of Models of Computation, Hong Kong, China, 20–22 May.

NIDISC: 16th International Workshop on Nature Inspired Distributed Computing, Boston, Massachusetts USA, 20–24 May.

UNCERTAIN REASONING: St. Pete Beach, Florida, USA, 22–24 May.

EI&I: Evolution, Intentionality and Information, University of Bristol, 29–31 May.

AIME: Artificial Intelligence in Medicine, Murcia, Spain, 29 May–1 June.

LoQI: Logic, Questions and Inquiry, Paris, France, 30 May–1 June.

GRADUATE EPISTEMOLOGY CONFERENCE: University of Edinburgh, 31 May–1 June.

JUNE

BSPS: British Society for the Philosophy of Science Annual Conference, University of Exeter, 4–5 June.

BAYSM: Bayesian Young Statistician Meeting, Milan, Italy, 5–6 June.

BISP: 8th workshop on Bayesian Inference in Stochastic Processes, Milan, Italy, 6–8 June.

LORI: 4th International Workshop on Logic, Rationality and Interaction, Hangzhou, China, 9–12 June.

CADE: 24th International Conference on Automated Deduction, Lake Placid, USA, 9–14 June.

NECESSITY, ANALYTICITY & A PRIORI: Oslo, 10–11 June.

ICAAIL: 14th International Conference on Artificial Intelligence & Law, Rome, Italy, 10–14 June.

INEM: Conference of the International Network for Economic Method, Erasmus University Rotterdam, The Netherlands, 13–15 June.

SocPHILPSYCH: 39th meeting of the Society for Philosophy and Psychology, Brown University, Providence, RI, 13–15 June.

TRoREC: The Reach of Radical Embodied or Enactive Cognition, University of Antwerp, 17–19 June.

LOGICA: Hejnice, Czech Republic, 17–21 June.

TAP: 7th International Conference on Tests and Proofs, Budapest, Hungary, 18–19 June.

GP@50: The Gettier Problem at 50, University of Edinburgh, 20–21 June.

ICFIE: 2nd International Conference on Fuzzy Information and Engineering, Kanyakumari, India, 22–23 June.

ISF: 33rd International Symposium on Forecasting, Seoul, Korea, 23–26 June.

HDIA: High-Dimensional Inference with Applications, University of Kent, Canterbury, 24–25 June.

CSR: 8th International Computer Science Symposium in Russia, Ekaterinburg, Russia, 25–29 June.

BW8: 8th Barcelona Workshop on Issues in the Theory of Reference, Barcelona, 26–28 June.

APPLIED PHILOSOPHY: Society for Applied Philosophy Annual Conference, University of Zurich, 28–30 June.

AIME: Artificial Intelligence in Medicine, Murcia, Spain, 29 May–1 June.

JULY

CAeITs: Causality and Experimentation in the Sciences, Paris, 1–3 July.

CEPE: Ambiguous Technologies: Philosophical Issues, Practical Solutions, Human Nature, Lisbon, Portugal, 1–3 July.

CIe: The Nature of Computation, Milan, Italy, 1–5 July.

ISIPTA: 8th International Symposium on Imprecise Probability: Theories and Applications, Compiègne, France, 2–5 July.

IC-EpsMsO: 5th International Conference on Experiments/Process/System Modeling/Simulation/Optimization, Athens, Greece, 3–6 July.

YSM: Young Statisticians’ Meeting, Imperial College London, 4–5 July.

CARNAP ON LOGIC: MCMP, Munich, 4–6 July.

ECSQARU: 12th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, Utrecht University, The Netherlands, 7–10 July.

AAP: Australasian Association of Philosophy Conference, University of Queensland, 7–12 July.

GDRR: 3rd Symposium on Games and Decisions in Reliability and Risk, County Cork, Ireland, 8–10 July.

ICALP: 40th International Colloquium on Automata, Languages and Programming, Riga, Latvia, 8–12 July.

SCEPTICISM: New Perspectives on External World Scepticism, MCMP, LMU Munich, 9–10 July.

IUKM: 3rd International Symposium on Integrated Uncertainty in Knowledge Modelling and Decision Making, Beijing, China, 12–14 July.

FoP: Foundations of Physics, LMU, Munich, 29–31 July.

AUGUST

WL4AI: Weighted Logics for AI workshop, Beijing, China, 3–5 August.

IJCAI: 23rd International Joint Conference on Artificial Intelligence, Beijing, China, 3–9 August.

WCP: 23rd World Congress of Philosophy, Athens, Greece, 4–10 August.

KSEM: International Conference on Knowledge Science, Engineering and Management, Dalian, China, 10–12 August.
LMoGDM: Logical Models of Group Decision Making, Düsseldorf, Germany, 12–16 August.
RACR: 4th International Conference on Risk Analysis and Crisis Response, Istanbul, Turkey, 27–29 August.
EPSA: European Philosophy of Science Association, University of Helsinki, Finland, 28–31 August.
EoM: Epistemology of Modality, University of Lisbon, 29–31 August.

SEPTEMBER

ICSCCW: 7th International Conference on Soft Computing, Computing with Words and Perceptions in System Analysis, Decision and Control, Izmir, Turkey, 2–3 September.
DiAL: Dialectic in Aristotle's Logic, Groningen, Netherlands, 2–4 September.
CSL: 22nd EACSL Annual Conference on Computer Science Logic, Turin, Italy, 2–5 September.
ECAL: 12th European Conference on Artificial Life, Taormina, Italy, 2–6 September.
ENPOSS: European Network for the Philosophy of the Social Sciences and the Philosophy of Social Science, University of Venice Ca' Foscari, 3–4 September.
MCU: Machines, Computations and Universality, University of Zurich, 9–12 September.
ITA: 5th International Conference on Internet Technologies and Applications, Glyndwr University, Wrexham, North Wales, UK, 10–13 September.
SUM: 7th International Conference on Scalable Uncertainty Management, Washington DC, 16–18 September.
CLPS: International Conference on Logic and Philosophy of Science, University of Ghent, 16–18 September.

PROGIC

The sixth workshop on Combining Probability and Logic.
 Special focus: combining probability and logic to solve philosophical problems. Munich, 17–18 September

TbiLLC: 10th International Tbilisi Symposium on Language, Logic and Computation, Georgia, 23–27 September.
AIAI: 9th IFIP International Conference on Artificial Intelligence Applications and Innovations, Paphos, Cyprus, 26–28 September.

COURSES AND PROGRAMMES

Courses

BFAS: Spring School on Belief Functions Theory and Applications, Carthage, Tunisia, 20–24 May.
NORDIC SPRING SCHOOL IN LOGIC: Nordfjordeid, Norway, 27–31 May.
ACAI SUMMER SCHOOL 2013: Computational Models of Argument, King's College London, UK, 1–5 July.
EASSS: 15th European Agent Systems Summer School, Kings College London, 1–5 July.
ESSLLI: 25th European Summer School in Logic, Language and Information, Heinrich Heine University in Düsseldorf, Germany, 5–16 August.
ETHICSCHOOL: Virtual Summerschool on Ethics of Emerging Technologies, 9–13 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.

PhD POSITION: in Logic and Cognitive Modelling, ILLC, University of Amsterdam, deadline 15 February.

PhD POSITION: in Philosophy, AOS: Analytic Philosophy / Logic / History and Philosophy of Science and Technology / Philosophy of Social Sciences / Philosophy of Mind and Cognitive Sciences, Tallinn University of Technology, Estonia, deadline 1 March.

PhD POSITIONS: in Philosophy of Science, University of Aberdeen, deadline 8 March.

PhD POSITIONS: in Science and Policy, Centre for Humanities Engaging Science and Society (CHESS), Durham University, deadline 11 March.

JOBS AND STUDENTSHIPS

Jobs

ASSISTANT PROFESSOR: in Logic or Analysis, Department of Mathematics, University of Connecticut, until filled.

POST-DOC POSITION: in Artificial Intelligence, Institute for Artificial Intelligence, University of Georgia, until filled.

POST-DOC POSITION: in Artificial Intelligence / Biomedical Informatics, Stevens Institute of Technology, until filled.

LECTURER: in Philosophy, AOS: Philosophy of Science / Mind / Language or Epistemology, University of Sussex, deadline 4 February.

POST-DOC POSITIONS: in Bayesian Inference, Department of Statistics, University of Oxford, deadline 8 February.

LECTURER: in Philosophy of Science, University of Oxford, deadline 8 February.

LECTURER: in Probability or Statistics, School of Mathematics, University of Bristol, deadline 11 February.

POST-DOC POSITIONS: in Philosophy of Social Science, TINT Centre of Excellence in the Philosophy of the Social Sciences, Helsinki, deadline 15 February.

POST-DOC POSITION: in Metaphysics of Science, Institut d'Histoire et de Philosophie des Sciences et des Techniques, Paris, deadline 15 February.

POST-DOC POSITION: in Philosophy and Cognitive Science, University of Murcia, Spain, deadline 15 February.

POST-DOC POSITIONS: in Philosophy and Science, The Rotman Institute of Philosophy, University of Western Ontario, Canada, deadline 15 February.

POST-DOC POSITION: in Philosophy of Science and Technology, Tallinn University of Technology, Estonia, deadline 1 March.

POST-DOC POSITION: in Statistics, University of Bristol, deadline 5 April.

POST-DOC POSITION: in Theoretical Philosophy working on "Infinite Regress" project, University of Groningen, The Netherlands, deadline 8 April.

Studentships

PhD POSITION: on project "Non-Classical Foundations of Mathematics," Department of Mathematics and Statistics, University of Canterbury, New Zealand, until filled.

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