The Reasoner

Volume 7, Number 12
December 2013

www.thereasoner.org
ISSN 1757-0522

Contents

Editorial

These days, it is pretty common that philosophers write joint papers with scientists. The interview with Vincenzo Crupi in the October edition of The Reasoner, for example, provides ample evidence for the interaction between reasoning researchers in psychology and formal epistemology. And of course, experimental philosophers have been blurring the boundary between philosophy and science since the turn of the century.

When asking myself when this trend started, it came to my mind that the pioneers of this approach might be found at Carnegie Mellon University (CMU). The philosophers at CMU have a long tradition in collaborating with researchers from other fields such as economics, psychology, statistics, computer science and machine learning. Examples include renowned scholars such as Teddy Seidenfeld, Peter Spirtes, Richard Scheines, and also rising young stars such as Kevin Zollman and David Danks. One name, however, may stand out in this illustrious company: Clark Glymour, founder of the department, Alumni University Professor at CMU, and perhaps one of the most fascinating and controversial scientists, methodologists and philosophers alive.

What makes Clark so fascinating is, inter alia, the fact that he started out as a philosopher of science, working on traditional topics such as a first-order logic formalization of the confirmation relation. His book *Theory and Evidence* (Princeton, 1980) sums up this early phase of his career. Over time, however, Clark's interests changed. After a detour through the history of science, in particular physics, he ended up in probabilistic modeling. Today, he might be best known for his work on Causal Bayesian Networks, in particular for his books *Discovering Causal Structure* (Academic Press, 1987, with Peter Spirtes, Richard Scheines and Kevin Kelly) and *Causation, Prediction and Search* (Springer, 1993, with Spirtes and Scheines). Looking at Clark's most recent output, one gains the impression that he has not fully abandoned philosophy of science, but that he spends most of his time working on scientific questions: searching for causal discovery algorithms, modeling interventions, solving forecasting problems, and so on. How come that a prominent philosopher of science makes such a radical shift? Reason enough to interview Clark Glymour for *The Reasoner*.

Jan Sprenger
TiLPS, Tilburg

Features

Interview with Clark Glymour

Jan Sprenger: Clark, clearly you are an intellectual multi-talent. What did attract you to philosophy in the first place, rather than pursuing a career in the sciences?

Clark Glymour: As a youth in summers free from high school drudgery I read Spinoza's *Ethics* and Darwin's *Origins*. I got more from Darwin. I tried putting the axioms and theorems and proofs of the *Ethics* into various mathematical notations—they pretended to be by parallel with Euclid after all—so I could follow them, but to no avail. (Not my fault—see what George Boole had to say on the topic.) At University I was captivated by the contrast between two of my philosophy professors, one a

Heideggerian, the other a student of Reichenbach's. In a way, I learned more from the Heideggerian, who in a seminar passed around a letter he had received from Martin, not for us to read (we were from Montana, so go figure) but to touch. I was nauseated, revulsed as I had been by the clergymen I had known in the small Catholic city in which I grew up. My distaste for awe is genuinely visceral. But I was inspired by Reichenbach's student, Cynthia Schuster, an older woman who had spent the second world war interned by the Vichy. She had stories. She knew stuff, and nothing awed her.

By the time the university expelled me because I refused to participate in required military classes, I had completed most of the requirements for a philosophy degree but I had two undergraduate years remaining. I enrolled in New Mexico where marching and saluting and assembling rifles were not required. I quickly discovered that much of the research produced by the philosophy department was utterly foolish, I mean utterly, utterly foolish—most of your readers will be too young or too wise to remember the metaphysical polarities and such of Archie Bahm, author of the *Directory of American Philosophers*. So what to do with two years? I thought about a second major in physics, but I had no mathematics, so I did a major in chemistry, where the mathematics was slow enough that I could catch up. Sometimes the race was close; I took quantum theory before I took calculus. Theoretical chemistry seemed a closed subject that was going to be governed by computational grinding, and in those days computing was excruciating; laboratory chemistry was for the agile and I am not. So evenings while the test tube racks bubbled away, I read *The Direction of Time* and decided to try for a Ph.D in History and Philosophy of Science. Schuster and Wes Salmon had been students together, so Wes took a gamble on me and I was admitted at Indiana, where I also continued to do courses in chemical physics.

So always, there were these two sides, which I thought of, and still do, as just one side: curiosity.

JS: How do you stand today toward your early philosophical works, e.g., the book *Theory and Evidence*? Should philosophy of science students still read it?

CG: The most important part of the book is the last paragraph. As to the rest, the problem was too big for me at the time I wrote it, and probably is still. I was only dimly aware of the literature in econometrics and engineering on parameter and system identifiability problems. I thought everything had to be done in full generality in first-order formalization or not at all, and I was not a good enough logician. I recently looked back at commentary on the book and found some of the commentary odd in a special way I did not appreciate at the time. But never mind that.

I have not for decades urged a student to read *Theory and Evidence*. I would (and do) urge them to think about the problem: how does evidence bear on hypotheses specifying properties and relations among variables that are not recorded in the data, and bear in such a way that truth or falsity of the hypotheses can eventually be determined? My work in later years caused me to recast the question as one of search or estimation. Teru Miyake is the only philosopher I know of (outside of CMU faculty and our former students) who has written anything sensible on the topic of search, a methodologically fascinating set of issues that have swept through the sciences in the last decades.

JS: What about your time as a farmer? How come that you temporarily withdrew from academia? And why did you come back again?

CG: Maybe the story has outrun the truth a little here. In academia, it doesn't take much to seem very unconventional. I had a sabbatical from Princeton, and spent it with my family in a hunters' cabin on the farm in Appalachia of an old friend (Don McCaig, still my oldest living friend) from Montana. We shared a milk cow and a litter of pigs and fun and worries and I wrote most of *Theory and Evidence* there. John Hicks, who worked at the Princeton accelerator, installed the cabin water pump, and David Malament came to visit.

In 1976 I left Princeton for the University of Oklahoma, not to become a farmer, but because my then wife deeply hated Princeton and wanted to move to her home state. But in Oklahoma, academia was pointless in more ways than I have space here to tell. I have some space. My first week in Oklahoma I went to the library—a handsome building of four stories—to get a copy of Hume's *Inquiry*. Couldn't find it in the card catalog so I asked and was told that *that sort of book* had to be got from interlibrary loan. What could be on the other three floors of the library? I went up to the 4th floor: empty; the 3rd: empty. But the 2nd: tons of outdated rations in case of nuclear war. So I rented a farm and raised some pigs, a calf, goats, chickens and whatnot, and I talked the library out of the expired rations and fed them to the pigs.

By 1978 I was divorced and other universities came calling with offers, and I accepted one from Chicago Circle. After an unhappy year there, I moved to Pittsburgh, where things have gone well.

JS: At some point you changed the focus of your research from classical topics in philosophy of science, including historical work, to causal modeling. How did this radical shift come about?

CG: Easy, just follow the problems and don't let stupid disciplinary borders stop you. *Theory and Evidence* had a chapter on "Causal Modeling" which I wrote after reading Hubert Blalock's little book on theories in the social sciences. Blalock had worked out what we would now call a "model selection" procedure for linear systems with Gaussian independent disturbances and exactly four variables. It was easy to see that he had done a very special case of a general problem that was at its base, graph theoretic, and if one could tell what constraints on observed variables a graphical model implied, that could be used as a search procedure. Fascinating. So I talked to every quantitative social scientist I could find and got nothing. I searched the literature for a month and found nothing (there was something, but I didn't find it at the time, 1980). I wrote to the most distinguished American graph theorist of the day, Frank Harary: no response. So frustrating, a question so key to so many fields, an answer that must be so close, and there was nothing except Blalock's little book and his four variable linear system. I couldn't let go. So I encouraged my students—Kevin Kelly, Peter Spirtes and Richard Scheines—to help me work on the problem. I bought a computer terminal for us to share (radical stuff at the time). We made a little progress generalizing Blalock's results to all acyclic graphs for linear systems with no unobserved variables. We did the same with quadratic correlation constraints for linear latent variable models. But no general understanding.

Then, in 1998, we read Pearl's book *Patterns of Plausible Reasoning*, and it became clear how to derive independence constraints from a graph. The pressing logical questions became computational. You could stop and say that is where philosophy ends, but it's not where curiosity ends. The space of graphical models increases superexponentially with the number of variables. You can't test every model, so how can you find the true one or any interesting causal properties of the true one? Spirtes and I worked literally day and night on the problem, more or less obsessed, getting more and more understanding of how models could be segregated by data, spurred by work of a colleague in computer science, Greg Cooper, who was working on the same problem. And then one afternoon Peter called me and said: "I think I see how we can do it." And that was the PC algorithm we published in 1991, followed by our book with Scheines, *Causation, Prediction and Search*. Those were the happiest intellectual years of my life. We had discovered something really important that no one else knew and done something that thousands of eminences said couldn't be done. I was really happy.

I was stunned at the response, and in some measure I am still bitter. A French probabilist, refereeing the manuscript for *Causation, Prediction and Search*, reported that the manuscript should be rejected because PC algorithm did not do what we said it did. He hadn't bothered to turn the page to read the second half of the algorithm where his objections were all met. David Freedman devoted essays and many lectures, including part of his Presidential address to the American Statistical Association, denouncing the book, deliberately misquoting. When I complained in a letter to him, he replied that, yes, he lied, so what? Persi Diaconis was little better, and reviews in the statistical literature by people who had obviously not read the book, were harsh. The statisticians at CMU were supportive or indifferent, and I am grateful for that. The book would never have been published without Stephen Fienberg's efforts with Springer.

The philosophical community has been disappointing in a different way. There were criticisms of assumptions, notably by Nancy Cartwright, which were honestly intended if not useful and not always well informed. I didn't mind that. I did mind the fellow who at a PSA meeting attacked *Causation, Prediction and Search* and, who, when his attributions were challenged by Kevin Kelly, replied that everyone knows the book is difficult so he hadn't actually read it. And I mind the repeated attribution by philosophers of our work to Pearl. Pearl did a great deal, he is a first-class, original and imaginative mind, but until our work he had thrice (like St. Peter) denied that graphical models could have causal significance, and his development of prediction algorithms derives from algorithms in *Causation, Prediction and Search*.

What the experience tells me is that disciplinary boundaries have become rocks in heads. The statisticians simply could not endure philosophers intruding in their domain; the philosophers simply could not acknowledge statistical work by philosophers as philosophy.

In October, CMU had a workshop on Causal Discovery in the Sciences. There were lectures by scientists on economics, ecology, gene regulation, climate modeling, cellular signaling, functional magnetic resonance imaging, clinical applications in autism research, and more. Much of the work—not all—involved developments from procedures in *Causation, Prediction and Search*, and some of it from applications I and my colleagues had inaugurated. In the scientific game, we have won: causality and automated search are legitimate ideas and methods with growing influence and application.

The Pittsburgh department of History and Philosophy of Science is a quarter of a mile away. To my knowledge, no faculty member or student from the department attended any workshop session. In the philosophical game, we have lost.

JS: Lots of your papers are co-authored with scientists. What is it that excites you about interdisciplinary work?

CG: Several things. We have learned a great deal about search algorithms by understanding what works and what fails in various scientific problems. My time of any original mathematics is past, but I can still do applications, and that is satisfying. I like seeing philosophical ideas—ok, some of *my* philosophical ideas—make a good difference in science. I have no laboratory—I don't even recognize most of the equipment in a modern biochemistry laboratory—but I take real pleasure in knowing that what I do know is lending a hand. By and large, the scientists are far more open than professional methodologists, i.e., statisticians. The scientists want to know how best to address very hard problems, and they are willing to listen and try strategies that make sense. Besides, if you are interested in, say, gene expression, or brain mechanisms, or whatever, the scientists are a lot more informative than the book reports that so often pass nowadays as philosophy of science.

JS: Let us also talk about your institution. Philosophy at Carnegie Mellon is well-known for its close links between philosophy and the sciences. Tell me why pursuing a Ph.D at Carnegie Mellon is commendable for students who want to do high-level philosophy. And what is your vision of philosophy education?

CG: I was asked in 1984 to buy CMU a philosophy department—it needed one in order to have a Phi Beta Kappa Chapter. I got the job partly because Dana Scott, a friend, and formerly a colleague at Princeton, was an influential figure at CMU, and because Jay Kadane in Statistics wanted someone who would hire his collaborator, Teddy Seidenfeld, which I was delighted to do. Basically, I had a free hand. The work on causal inference was on the verge of breakthroughs, I knew it, and I needed resources to keep my then graduate student collaborators, so I took the job and hired them as assistant professors.

The department has one unique qualification for appointments: research that has a substantive bearing on mathematics, empirical science, or public policy. All of my colleagues have and do, and we expect our graduate students to do the same. So, if a student is motivated by philosophy of science but wants to do something that matters in these domains, Carnegie Mellon is the place to come if one can. If that isn't the ambition, it's the wrong place.

I think the philosophy of science education received by students in most departments in the United States and Britain is appalling; the way it is appalling is the standards, the expectations, and the skills. My evidence is from HPS at Pitt, and from the many job applications I review. Basically, the students learn a standard literature—Hempel, Kuhn, van Fraassen—and something about some science about which they write a book report with commentary, called a doctoral thesis. Of course, there are exceptions, good exceptions, but I don't see many. In general, the students can't write an algorithm, can't code, and can't do any but the most elementary statistics or probability theory, and even the mathematically talented ones are often terrified that actually proving something formal and scientifically relevant will cost them a career—their faculty advisors warn them as much. What they learn in graduate school is that high-minded book reports are what they should do, and by the time they receive their doctorates, it is all that they can do.

As to philosophy more broadly, the New York Times has just run an article on the fading of humanities at Stanford, taken as a bell-weather. Disciplines die hard, but some disciplines, or aspects of them, should be left to wither and die as research centers that universities support. English departments, for example, and a good part of philosophy, which has become so absorbed with its own history that it bids fair to swallow itself. We need to teach the history of philosophy; we don't need yet another book on Kant.

JS: The most famous quote about you is probably by Richard Jeffrey: "Who cares whether a pig farmer is a Bayesian?". It refers to your temporary timeout from academia and your essay "Why I am not a Bayesian". After thirty-three years, have your standpoints toward Bayesian modeling in philosophy changed?

CG: No, although the set of reasons I would give today are not exactly those I gave in "Why I am Not a Bayesian", but some of them are. I think Bayesian epistemology evades lots of hard questions and provides few insights except internally, among the Bayesians to the Bayesians. Bayesian epistemology is too easy, with endless free parameters for any problem. Take one problem: computability. Computability limits ability quite as much as does the finiteness of our evidence, and Bayesian epistemology is intractable. Simple example: Suppose you have a probability measure on sentences of a first order language, as in Gaifman and Snir's famous paper (at least it should be famous). Then by Turing's theorem you have to give 0 probability to an uncomputable set of contingent sentences. Why should that be a rational ideal? Take another: We have the classical convergence theorems but little attention to their limitations or how little they guarantee. There is a nice essay forthcoming on the issue in *Philosophy of Science*. The endless literature on how properly to define a Bayesian confirmation function has no target, no practice it represents or plausibly idealizes. Take the probabilistic accounts of explanatory power, none of which (except one, by David Glass) are published with considerations of their accuracy in hypothesis selection, or their adequacy as theories of psychological judgement (one exception, a paper by Jonah Schupbach, which is serious and well intended and ruined by incompetent data analysis. More evidence that young philosophers of science are not taught how to do science.)

These complaints are about Bayesian epistemology, not Bayesian statistical methods, which I often use. I just don't buy or even think intellectually serious the rhetoric around them.

JS: One might expect that somebody who is very much into causal and statistical reasoning has also a lot to say on the foundations/philosophy of statistical inference. How come that you are relatively silent on this topic?

CG: I said most of what I have to say on frequentist versus Bayesian foundations for statistics in one essay ("Instrumental Probability," *The Monist*, 2001). I have written sections of several books and essays on issues now gaining attention, for example Matt Kotzen's essay in *Nous* (2001) largely sharing (but not citing) my views on multiple hypothesis testing. But I never wrote an essay with, say, the title "Multiple Hypothesis Testing" or "The Idiocy of Bonferroni Adjustments." So I have written about the subject, but it's a question of how big that section is in my market shelf, and it's small and in corners where not many people go.

JS: How do you look upon the state of philosophy, and philosophy of science in particular, as an academic discipline? Where do you see its future in?

CG: Look at recent books by senior American figures. Kyle Stanford's *Exceeding Our Grasp* is an elaboration of a tautology: if the truth is something you haven't thought of, then you won't think the truth. Paul Churchland's *Plato's Camera* is a summation of his thought that becomes a combination of pseudo neuropsychology and conventional platitudes whenever he attends to scientific method (see his amazing discussion of Darwin). Sandra Mitchell's *Unsimple Truths* is a well intentioned collection of remarks that are utterly banal to most scientists, and should be to philosophers. Bill Harper's *Isaac Newton's Scientific Method* is a loving explication of Newton's argument in the third book of the Principia, but as to Newton's method, Harper says it's not bootstrapping, but its sort of *like* bootstrapping and its really wonderful. For what it is, in a concrete way, a method someone could follow, we will have to wait.

In some respects, things are better in Europe. Formal methods are used ingeniously and rigorously to address general questions, but for the most part not questions with much, or any, bearing on science. How do I use a ranking function with noisy sample data to find out anything? Franz Huber wrote me a sketch, but with no example; Wolfgang Spohn told me he had no idea. Maybe someone will.

As to the future of the subject, the past is prologue.

JS: Finally, I would like you to give the opportunity to make a point that you find important, but that has not been covered so far in the interview. Any thoughts?

CG: Yes, thanks for the opportunity and your flattering introduction to the questionnaire.

Mizrahi's argument against Phenomenal Conservatism

Mizrahi (2013, "Against Phenomenal Conservatism", *The Reasoner*, 7(10), pp. 117–118) argues that Phenomenal Conservatism (see Huemer 2007, "Compassionate Phenomenal Conservatism", *Philosophy and Phenomenological Research*, 74, pp. 30–55) is an untrustworthy method of fixing belief (MFB). I respond that Mizrahi's argument is unsound because one premise is rationally unacceptable, and that if this premise is refined and made more acceptable, the argument proves invalid.

Phenomenal Conservatism says that *seemings* are special mental states—i.e., propositional attitudes, different from beliefs, capable of supplying justification for their contents. Accordingly:

> (PC) If it seems to S that p, then, in the absence of defeaters, S thereby has at least some degree of justification for believing p.

(PC) holds that it is *by virtue of* (or *on the grounds of*) S's having a seeming with content p that S has some degree of defeasible justification for believing p (cf. Huemer 2007, p. 30).

This is Mizrahi's argument:

1. (PC) [Assumption for reductio]

2. It seems to S_1 that *p* and it seems to S_2 that not-*p*, independently of each other. [Premise]

3. Therefore, in the absence of defeaters, S_1 has some degree of justification for believing *p* and S_2 has some degree of justification for believing not-*p*. [From (1) & (2)]

4. If an MFB provides some degree of justification for contradictory beliefs, it's untrustworthy. [Premise]

5. Appealing to seemings provides some degree of justification for contradictory beliefs. [From (3)]

6. Therefore, appealing to seemings is an untrustworthy MFB. [From (4) & (5)]

As Mizrahi indicates, (2) appears true for some *p*. For instance, to Jackson (1982, "Epiphenomenal Qualia", *Philosophical Quarterly*, 32, pp. 127–136) it *seems* that Mary learns something new, whereas to Dennett (1991, *Consciousness Explained*, Boston: Little Brown) it *seems* that she doesn't. To Hauser (2002, "Nixin' Goes to China", in Preston and Bishop (eds.), *Views Into the Chinese Room*, NY : OUP) it *seems* that the person in the Chinese room understands Chinese, whereas to Searle (1999, "The Chinese Room", in Wilson and Keil (eds.), *The MIT Encyclopedia of the Cognitive Sciences*, MIT Press) it *seems* that that person doesn't. I find Mizrahi's examples prima facie plausible, so I won't question (2).

Mizrahi reports an objection to (4) by an anonymous reviewer, which he leaves unaddressed but appears to regard as serious. I don't think Mizrahi's argument is flawed because of it. The objection runs as follows: you know that an urn U contains a red, a blue and a yellow ball. Alice extracts one ball from U but you cannot see its colour. She truthfully tells you that (*e*) the ball isn't yellow. This gives you some justification to believe that (*r*) it is red and some justification to believe that (*b*) it is blue. The alleged difficulty for (4) is that although *r* and *b* are incompatible, your MFB isn't untrustworthy. I see no real challenge for (4) because *r* and *b* are *incompatible* but not just one the *logical negation* of the other. (It is customarily accepted, for instance in science, that the same evidence can support *incompatible* hypotheses.) Furthermore, in this example it is false that *e* gives you some justification for *both* *r* and not-*r* (or both *b* and not-*b*). Suppose U contains the three coloured balls only. It is intuitive that before you learn *e*, your degree of confidence in not-*r* should be $2/3$, but after you learn *e* your confidence in not-*r* should *drop* to $1/2$. So *e* cannot give you justification for not-*r*. This conclusion holds in general even if U contains *additional* balls. Since *e* boosts your confidence in *r*, it follows from the probability calculus that *e* must *lessen* your confidence in not-*r*.

Mizrahi's argument against (PC) presupposes that the same MFB can be used by different subjects (or the same subject at different times). But there is a problem with the way he implements this idea. In particular, if (4) were true, we should conclude that *any* MFB utilizable by different subjects is untrustworthy. For, trivially, different subjects may have different evidential grounds that support contradictory propositions via the same MFB. Take for instance testimony. Let p be the proposition that my pet flies. I tell S_1 that my pet is a bird and I tell S_2 that my pet is a penguin. S_1 will have some justification for believing p and S_2 some justification for believing not-p. So if (4) is true, testimony is untrustworthy. Consider now perception. Suppose S_1 only sees that my pet has a beak, whereas S_2 clearly sees that my pet is a penguin. S_1 has some justification for believing p and S_2 has some justification for believing not-p. If (4) is true, perception is untrustworthy. These examples easily multiply.

We cannot accept (4) because this would commit us to a very implausible conclusion. An obvious refinement of (4), which settles this difficulty, is the following:

> 4*. If an MFB provides some degree of justification for contradictory beliefs *on the grounds of the same evidence*, it's untrustworthy.

Mizrahi might intend (4) as equivalent to (4*). For instance, to defend (4) Mizrahi envisages a situation in which he uses a Litmus test as a MFB about the pH of a given solution. The test is repeated again and again. Mizrahi sensibly concludes that if his blue Litmus paper sometimes turned red (thereby indicating an acidic solution) and sometimes stayed blue (thereby indicating a basic solution), he wouldn't put much trust in his MFB. In this thought experiment, Mizrahi's MFB can be described as processing at different times *the same evidence*, constituted by the same solution and the same background information necessary to interpret the test's observational outcomes.

Suppose we replace (4) with (4*). The resulting variant of Mizrahi's argument against (PC) would go through only if (5) could be interpreted accordingly, i.e., as stating that appealing to seemings provides some degree of justification for contradictory beliefs *on the basis of the same evidence*. This interpretation is very questionable. Take again the case in which it seems to Jackson that (p) Mary learns something new, whereas it seems to Dennett that she doesn't. The phenomenal conservative would claim that the *evidential grounds* of Jackson's belief that p and the *evidential grounds* of Dennett's belief that not-p are to be identified with Jackson's and Dennett's respective seemings. The phenomenal conservative would thus insist that Jackson and Dennett have (defeasible) justification for contradictory beliefs because they have *different* evidence constituted by their conflicting mental states coinciding with incompatible seemings. This example generalizes: appealing to seemings can provide some justification for contradictory beliefs only on the basis of *different* evidential grounds—i.e., different seemings. In conclusion, if we replace (4) with (4*), (6) doesn't follow from (4*) & (5); the resulting argument is invalid. I don't exclude that Mizrahi is onto something and that (PC) could turn out to be untrustworthy. However, to believe so we would need a neat argument that Mizrahi has not delivered.

Luca Moretti
Philosophy, University of Aberdeen

'This sentence'? Which sentence?

Samuel Alexander is slightly doubtful in his concluding sentence in (2013: "This sentence does not contain the symbol X", *The Reasoner* 7.9, p. 108). He concludes: 'the

English sentence "This sentence does not contain the symbol X" certainly does appear to contain the symbol X'. Whatever could have caused his lack of confidence about the obvious? The sentence 'This sentence does not contain the symbol X' certainly *does* contain the symbol X!

Showing doubt at this final point, though, is not the only uncertainty Alexander displays, and at an earlier point it is far more appropriate. For he says about his constructed fixed point λ that 'we feel tempted to gloss λ as "This sentence does not contain the symbol X"'. This temptation is what should be resisted. For λ *cannot* be glossed in the way Alexander proposes. He is working 'in the language of Peano Arithmetic extended with a new symbol X', and that language does not contain any demonstratives. 'This sentence' is a demonstrative expression.

Getting clear that λ does not say what Alexander thinks it does will not clear up the whole problem he has, but it is a start. Alexander produces a formula ψ (in the language of Peano Arithmetic without X) which holds of the Gödel number of a formula ϕ just so long as X occurs in ϕ. He then uses the fixed-point theorem to generate a sentence λ (not containing X) such that λ is logically equivalent to 'λ does not contain X'. He concludes (quite rightly) that λ is true. But equivalence is not identity: 'it is triangular' is equivalent to 'it is trilinear', for instance, in the right setting, yet triangularity is not trilinearity. So propositional identity is given by translations and synonymy, not logical equivalence, and that means that λ does not say anything about the sentence λ itself, even if it is equivalent to a proposition about that sentence. This fact is no more surprising than having a sentence that does not contain the symbol X being equivalent to a sentence containing the symbol X (as in the fixed-point theorem just mentioned).

So it is a common misunderstanding of sentences like 'This sentence does not contain the symbol X' that is causing the trouble that Alexander has. This further point is perhaps best seen in connection with the simpler sentence that is sometimes said to generate the Truth Teller Paradox: 'This sentence is true'. The most obvious feature of this sentence is that it contains a demonstrative 'this sentence', and so in particular Tarski's unamended Truth scheme does not apply to it. That is to say, one cannot say

> 'This sentence is true' is true if and only if this sentence is true,

any more than one can say

> 'he is happy' is true if and only if he is happy.

If truth is still to be attached to such sentences then the Tarskian Truth scheme must be modified in some way:

> 'he is happy' said of John is true if and only if John is happy,
>
> 'this sentence is true' said of s is true if and only if s is true.

In this way it becomes apparent that the sentence 'this sentence is true' has many different uses, and so can be used to express many different propositions. For instance, it can be spoken out loud accompanied with a gesture to a sentence written on a blackboard, and it can be written with the sentence referred to exhibited after a colon in a following clause. Thus one may express the proposition that 'the earth is round' is true this way:

This sentence is true: ‘The earth is round’.

In the absence of any such attached further sentence, of course, the most salient object for the demonstrative ‘this sentence’ to be taken to refer to is the very sentence itself, in which case the sentence becomes what some might want to call ‘self-referential’. But the sentence is not then self-referential in itself, since its subject phrase has to be given the specific ‘self-referential’ interpretation in order for this to happen. And the most important thing to notice then is that the specific interpretation then involved makes its subject phrase refer to a sentence that contains a subject phrase without a determinate referent. For in

This sentence is true: ‘This sentence is true’,

the first ‘this sentence’ has a reference, but the second ‘this sentence’ has none, because at that second occurrence it is within quotation marks, and so is only mentioned, and not used. There is a proposition expressed in the ‘self-referential’ case, and it is simply that ‘This sentence is true’, i.e., the sentence before an interpretation is given to its subject phrase, is true. But the sentence itself can have no truth-value, entirely because its subject phrase, being uninterpreted, has no determinate referent. So the proposition then expressed is simply false, while the associated sentence is neither true nor false.

In a like fashion readers can now be more confident that the sentence ‘This sentence does not contain the symbol X’ (in which, as should be clear after the above, no specific reference is given to ‘this sentence’) does contain the symbol X. Thinking otherwise involves a use-mention confusion.

HARTLEY SLATER
Philosophy, UWA

NEWS

Announcement: *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence*

Ray Solomonoff (1926–2009) considered the notion of a Universal Turing Machine (UTM), a simple abstract platonic mathematical model with—as far as we know, according to the Church-Turing thesis—all the power of today’s computers. In similar spirit to Ockham’s razor, a body of data which can be generated from a short input to a UTM can be considered to be simple in a precise sense. With much mathematical sophistication, in the 1960s Solomonoff developed this to arrive at a notion which he called “algorithmic probability”—essentially the a priori probability of a data-set. For the mathematicians and statisticians, Solomonoff contended that this gave optimal Bayesian prediction, and (in 1978) proved powerful convergence results. One of the many repercussions of this work which Solomonoff was aware of was and is its influences in philosophy of science and language.

The Solomonoff memorial conference proceedings *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence* (Springer LNAI/LNCS 7070) includes articles using the above-mentioned approach to discuss philosophical matters including but not limited to quantification of simplicity, quantification of surprise and unexpectedness, quantifying originality or creativity, falsification, algorithmic metaphysics, Ockham's razor and Goodman's grue paradox.

David L. Dowe
Information Technology, Monash

Logic and Philosophy of Science, 16–18 September

From Monday 16 to Wednesday 18 September 2013, the Centre for Logic and Philosophy of Science brought philosophers and scientists together at Ghent University (Belgium) to discuss various issues pertaining to logic and philosophy of science. The conference marked the 20th anniversary of the centre, which, from its foundation in 1993, maintains a research tradition focused on the logical, methodological and epistemological analyses of scientific reasoning processes. The conference was opened by the centre's founder, Diderik Batens. He argued that adaptive logics provide us with ways to formalise defeasible methods and their combinations. He presented an extensive overview of tools that combine adaptive logics in standard format, with special focus on combinations that formalise prioritisation. Philosophers and scientists can choose one of these specific formal tools depending on their applications. Batens concluded that, because many of the presented solutions for complex combinations have a simple dynamic proof theory, they are well suited to formalising actual reasoning processes. Next, Natasha Alechina reconsidered the logical omniscience problem, arguing that it is particularly pressing when formally modeling the reasoning of resource-bounded agents. After presenting a historic overview of various accounts of such reasoning, she explored the possibility of interpreting inferences as a specific type of actions, and applying a specific type of action logics to these. This results in a rich formal system, which allows one to represent the time, resources and communication needed for agents to arrive at a given piece of knowledge explicitly in the object language.

Tuesday opened with a lecture by Hanne Andersen, who challenged the Kuhnian idea of an autonomous agent, arguing that scientific practices are governed by cross-disciplinary activities and subspecialisation. She devised a framework to better understand science by pinpointing three interlocking continuums that focus on actual research activities: cognitive convergence versus divergence, epistemic dependence versus independence, and hierarchically enforced versus shared cooperative cognition. She concluded that research can best be understood as an interplay between contributory and interlocking expertise. This was followed by another plenary lecture, by Graham Priest, who showed how logic can be rationally revised by assessing a logic as a mathematical theory against more pre-theoretic norms of logical correctness present in natural language, employing standard criteria of theory choice. The final lecture of the day was given by Hasok Chang, who offered three arguments against the possibility of chemistry being reduced to physics. First, the foundation of quantum chemistry is in fact classical and rooted in 19th century organic structural chemistry. Second, rather than deducing chemical theories from physics, chemists help themselves to the conceptual resources of physics as the need arises. Finally, physics itself is hardly the unified science reductionists often assume it to be. Chang concluded that together, these considerations make the prospect of reducing chemistry to physics seem bleak.

The final day was opened by Stephen Read. In reaction to Arthur Prior's tonkish attack on the proof-theoretic justification of logical laws, Michael Dummett introduced the notion of harmony between elimination-rules and introduction-rules: rules for eliminating and introducing formulas in a proof should 'match' one another in the sense that elimination-rules add no more and no less to whatever meaning is encapsulated in the introduction-rules. Read showed how for classical logic higher-level rules, i.e., rules assuming as premises a derivation rather than a formula, can be 'flattened'. These rules can be reduced to stable rules that discharge only formulas, not derivations. The result no longer holds in a constructivist setting. Finally, James Woodward explored the scope and limitations of the mechanistic approach to explanation in biology from an interventionist perspective. Although the mechanistic approach is a useful way to understand biological explanation, he took issue with the view that mechanistic models that provide more lower-level detail are always explanatorily superior to those that do not. Instead, he argued for a contrastive view, according to which many successful explanatory models abstract away from lower-level detail.

The contributed sessions were divided into general sessions and specialized symposia. The general sessions covered logic (Demey, Więckowski, de Araújo, Ficara, de Vos, Luczak, Lutz, D'Alfonso) and philosophy of science (Scorzato, Feest, Wenmackers, Gauderis, Kao, de Winter, Kosolosky, Heesen). Specialized symposia were held about mathematics and computation (Daylight, Riss, Allo, Van Bendegem, van Kerkhove), causality in medicine (Bulcock, Osimani, Krueger), logic and action (Sergot, Kulicki, Trypuz), proofs, programs and procedures (Dechesne, Kramer, Jespersen), functional analysis and explanation (Wouters, Vermaas, Huber), paradoxes in non-classical logic (Verdée, Van De Putte, Raclavský, Angelova, Başkent, Andreas, Carnielli, Rodrigues Filho and Omori), and rational disagreement in science (Šešelja, Straßer, Leuschner, Biddle, Seidel, Kruse, Wieland, van der Kolk).

The conference closed with a celebratory reception. The Centre hopes to continue its research in logic and philosophy of science for many years to come.

Raoul Gervais
Laszlo Kosolosky
Frederik Van De Putte
Mathieu Beirlaen
Rafal Urbaniak
Peter Verdée
Centre for Logic and Philosophy of Science, Ghent University

Combining Probability and Logic, 17–18 September

The Sixth Workshop on Combining Probability and Logic (Progic 2013) continued the biannual Progic workshop series. It was organized by the Progic steering committee consisting of Jeff Helzner (Department Philosophy, Columbia University), Niki Pfeifer

(Munich Center for Mathematical Philosophy, LMU Munich) Jan-Willem Romeijn (Department of Philosophy, University of Groningen), Gregory Wheeler (Centre for Research in Artificial Intelligence, New University of Lisbon), and Jon Williamson (Department of Philosophy & Centre for Reasoning, University of Kent). Progic 2013 took place at the Carl Friedrich von Siemens Stiftung (Nymphenburg Palace in Munich, Germany) on September 17 and 18, 2013. Niki Pfeifer served as the local organizer.

Five invited talks and nine contributed talks delivered by philosophers from various parts of the world focused on the Progic 2013 theme "*Combining probability and logic to solve philosophical problems*". Each submitted contribution was marked blindly and the top 20% were accepted to be presented at the workshop: The high number of mostly high quality submissions is a strong indicator of the vitality of the field.

Igor Douven (Faculty of Philosophy, University of Groningen), Alan Hájek (School of Philosophy, Australian National University), Kevin T. Kelly (Department of Philosophy, Carnegie Mellon University), Hannes Leitgeb (Munich Center for Mathematical Philosophy, LMU Munich), and Peter Milne (Department of Philosophy, University of Stirling) were invited speakers. Douven's talk *Conditionals and inferential connections* is based on joint work with Shira Elqayam, David Over, Henrik Singmann, and Janneke van Wijnbergen-Huitink. Douven argued that truth conditions of conditionals should be constituted by inferential connections and he supported this claim by experimental data. Hájek's talk *Probabilities of counterfactuals and counterfactual probabilities* defended the thesis that counterfactuals have truth conditions. Kelly's talk (joint work with Hanti Lin) *Acceptance without certainty or stability* contained new results on how Bayesian credal states map propositional belief states and on how qualitative reasoning tracks Jeffrey conditioning. Leitgeb asked in his talk *The Humean thesis on belief: Belief and stable probability* how rational belief relates to degrees of belief and concluded that the former corresponds to resiliently high probability. Milne's talk *Information, confirmation, and conditionals* defined information-added conditionals and discussed their properties.

Glauber De Bona (Department of Computer Science, University of São Paulo), Liam Bright (Department of Philosophy, Carnegie Mellon University), Hykel Hosni (London School of Economics), Teddy Groves (Philosophy Department, University of Kent), Jürgen Landes and Jon Williamson (both Philosophy Department, University of Kent), Arthur Paul Pedersen (Max Planck Institute for Human Development, Berlin), Dana Scott (Department of Mathematics, University of California, Berkeley), Stanislav O. Speranski (Sobolev Institute of Mathematics, Novosibirsk), and Sean Walsh (Department of Logic and Philosophy of Science, University of California, Irvine) presented contributed talks. De Bona's talk (joint work with Fabio G. Cozman and Marcelo Finger) discussed nesting of probabilistic operators and probabilistic expressivity as criteria for classifying propositional probabilistic logics. Bright's talk investigated how degrees of incoherence can be measured. Hosni (joint work with Tommaso Flaminio and Lluís Godo) presented a logical analysis of de Finetti's notion of event within the framework of information frames. Groves argued that Carnapian inductive logic should be preferred over falsificationism in the philosophy of statistics. Landes and Williamson presented joint work which aims to find a new and unified justification for objective Bayesianism. Pedersen explained a key lemma behind a full numerical representation of strictly coherent, possibly non-Archimedean, preferences in terms of subjective expected utilities formed from possibly non-Archimedean probabilities and utilities. Scott presented work in progress on developing a stochastic λ-calculus. Speranski discussed expressibility and computability aspects of quantified probability logics, and connections between such logics and elementary analysis. Walsh defended the empiristic thesis that arithmetical knowledge may be extended by probabilistic confirmation in the same sense as it may be by proof.

The workshop was generously supported by the Alexander von Humboldt Foundation (Munich Center for Mathematical Philosophy) and by the Carl Friedrich von Siemens Stiftung.

Selected papers of Progic 2013 will appear as a special issue in the Journal of Applied Logic. Program, abstracts, and videos of the talks are available at the workshop website.

The next Progic workshop will focus on *formal epistemology and inductive logic* and will take place at the University of Kent (Canterbury, UK) in April 2015. It will be preceded by a Spring School on Combining Probability and Logic.

Papers from the previous Progic workshop, held at Columbia University, have been published this month in a special issue of the Journal of Applied Logic, edited by Jeff Helzner.

Niki Pfeifer
Munich Center for Mathematical Philosophy,
LMU Munich

Annual Buffalo Experimental Philosophy Conference, 11–12 October

The field of experimental philosophy is rapidly expanding its subject matter, improving existing methodologies, and introducing new techniques to improve analysis. The University at Buffalo has fostered these innovations, hosting three conferences dedicated to experimental philosophy since 2009. Organized by Associate Professor of Philosophy James Beebe and UB philosophy PhD candidates J. Neil Otte and Paul Poenicke, this year's conference reflected the growing maturity of the discipline with thirteen presentations over two days, each sharing experimental data. Speakers received detailed feedback from leading experimental philosophers in an informal, workshop environment.

Anthony Jack opened the conference, presenting neurological evidence suggesting that underlying structures in the brain give rise to intuitions in support of dualism. Nicole Hassoun presented evidence gathered from Kiva, an online, non-profit organization that allows individuals to aid other individuals around the world. She argued her data suggests we should incorporate both a threshold and a prioritarian principle into the analysis of what principles for aid distribution people accept. Peter Blouw, Wesley Buckwalter, and John Turri reported data showing knowledge attributions are highly sensitive to lucky events that change the explanation for why a belief is true. By contrast, they argue, knowledge attributions are surprisingly insensitive to lucky events that threaten but fail to change the explanation for why a belief is true.

In 2004 Edouard Machery, Ron Mallon, Shaun Nichols, and Stephen Stich published what has become one of the most widely discussed papers in experimental philosophy, in which they reported that East Asian and Western participants had different intuitions about the semantic reference of proper names. A flurry of criticisms of their work has emerged, and although various replications have been performed, many critics remain unconvinced. James Beebe presented a review of the current debate over Machery et al.'s (2004) results and then reported the results of studies that reveal significant cross-cultural and intra-cultural differences in semantic intuitions when controlling for variables that critics allege have had a potentially distorting effect on Machery et al.'s findings. He argued these results confirm the robustness of the cross-cultural differences observed by Machery et al. and thereby strengthen the philosophical challenge they pose.

Edouard Machery provided the keynote presentation for the conference. He argued that philosophers and ordinary people do not conceive of subjective experience in the same way. He provided experimental support for this hypothesis before proposing that for the folk, subjective experience is closely linked to valence.

The conference was highlighted by the variety of talks, as other speakers reviewed findings on philosophical expertise (Jennifer Nado), gender (Garrett Marks-Wilt), and folk mereology (David Rose). Innovative research on epistemic egocentrism (Joshua Alexander, Chad Gonnerman, John Waterman) and the Knobe Effect (Brian Robinson) advanced the systematic exploration of these subjects acutely examined in the literature. With a conference scheduled for 2014, the University at Buffalo is set to host another intense, synoptic review of experimental philosophy.

James Beebe
J. Neil Otte
Paul Poenicke
Philosophy, University at Buffalo

Probabilistic Modeling in Science and Philosophy, 11–12 October

Probabilistic models such as random walk or percolation models are successfully used in many disciplines of the natural and social sciences. Despite their frequent uses, these models have not received much attention by philosophers working on either probabilities or modeling. To initiate a methodological and philosophical discussion about probabilistic models, Claus Beisbart and Christoph Raible (Bern) organized a conference, which took place at the Oeschger Centre for Climate Change Research of the University of Bern on October 11–12, 2013. It brought together practitioners and philosophers; a special emphasis was put on modeling in climate science.

The workshop started out with talks that reviewed the state of the art of probabilistic modeling. Stephan Hartmann (Munich) showed how probabilistic models can be used to address philosophical questions. As one example, he presented a theoretically motivated model of joint deliberation. Christoph Raible (Bern) discussed the current practice of climate scientists to use model ensembles to handle various sorts of uncertainties. Dirk Helbing (Zurich) drew attention to risks that arise in complicated systems with a huge number of internal connections and advertized what he calls a 'global system science'. Finally, Frank Schweitzer (Zurich) reviewed the conception of Brownian agents. These are building blocks that constitute the system under consideration; their mutual interactions are described using probabilities.

A second group of talks took a more theoretical perspective and discussed the scope and limitations of probabilistic models as well as reasons that may motivate the use of probabilistic models. Wendy Parker (Durham) focussed on a method called stochastic parametrization currently debated in climate science. The idea is to account for sub-grid processes not resolved in a climate model by using probabilities. Margaret Morrison (Toronto) compared three methods to quantify uncertainties in validation, viz., validation metrics, hypothesis testing and Bayesian probabilities. Her conclusion was that Bayesian probabilities are most suited to handle uncertainties in model validation.

Since probabilistic models are often used for forecasts, Johanna Ziegel (Bern) defined a statistical framework for probabilistic forecasting. She particularly focused on the evaluation of probabilistic forecasts. Roman Frigg (London) discussed the prospects of probabilistic forecasts under the conditions of uncertainties. He showed that probabilistic forecasts are sensible if there are only uncertainties about the initial conditions. Using a concrete model of population dynamics, he argued that probabilistic forecasts fail if model uncertainty is present. Seamus Bradley (Munich) took up the question of how we can more successfully deal with model uncertainty. He defined non-probabilistic odds to extract at least some useful information from models that are affected from this type of uncertainty.

The last few talks addressed philosophical questions of how probabilistic models represent and what the related probabilities mean. Claus Beisbart (Bern) analyzed probabilistic models as inferences and argued that we can often usefully distinguish between two sorts of probabilities in modeling, viz., probabilities that describe the source and probabilities that apply to the target. He developed a view according to which both types of probabilities are Bayesian, but left open the possibility of other interpretations. Aidan Lyon (Maryland) raised the question of how well personal confidence is measured using probabilities. According to empirical findings, people tend to be over-confident if their confidence is measured using a single probability. Lyon argued that this problem is alleviated if confidence is quantified using intervals of probabilities. Finally, Rafaela Hillerbrand (Delft) discussed probabilistic climate models against the background of reductionism. She argued that less detail can sometimes be more.

The lively discussions at the workshop showed that probabilistic models are a rich topic with many connections to other philosophical issues. Abstracts of the papers are available here.

Claus Beisbart
Philosophy, Bern

Expressing discontent: appropriate or not? And if so, when, where, and how? 25 October

The 'Philosophical Activism' initiative—organized by the Centre for Logic and Philosophy of Science (CLWF) and the Centre for Ethics and Value Inquiry (CEVI) at Ghent University (Belgium)—comprises a series of one-day workshops that focus on the depths and widths of what it means to be philosophically active. In our second workshop on October 25, 2013, we welcomed papers that

> (1) examined what the notion of discontent might entail from a philosophical perspective, and
>
> (2) elaborated on how discontent can and should be (philosophically) expressed on the border between science and society.

Together with our 6 contributed speakers and audience we tackled the overarching question to what extent civil and/or scientific discontent can have their place in democracy/science?

Machteld Geuskens (Tilburg University) tried to identify a few key factors philosophical activists could and should have. An activist role for political philosophers is, according to her, to mark and defend the difference between activists and terrorists, and to publicly question the reactions of those in power.

Bob Brecher (University of Brighton) put forth arguments as to why academic work does not preclude activism and often requires it, hinting that disinterest does not necessarily entail being uninterested. Academics thus have a particular responsibility to act on their knowledge: to act against their knowledge is a form of corruption.

Mathijs van de Sande (KU Leuven) conceptualised how our discontent with the 'status-quo' can lead us to a critical but realist stance towards potentials for—and forms of—radical change. An 'optimism of the intellect and pessimism of the will' should be endorsed as an attitude that allows us both to politically support and scientifically understand and recognize such forms and potentials.

Joris Luyckx (Ghent University) offered arguments for an alternative ethics, called prefigurative ethics, as opposed to deontological and teleological alternatives. This ethics would serve as a means to bridge the ever widening gap between doing (i.e., expressing-out-loud) and thinking (i.e., disapproving-in-theory) of contemporary academia.

Wim Vandekerckhove (University of Greenwich) elaborated on his extensive work related to whistleblowing in and by organisations. In this talk he explained why the wrongdoing starts once recipients of a concern respond to whistleblowers. Wrongdoing in an organization should not necessarily worry the wider public, but wrongdoing by the organization, so he argued, is always in the public interest.

Anna de Bruyckere (Durham University) showed by using an example from sexology in 1970s America what expressing lay discontent does not necessarily entail: the possibility of escaping entrenched discursive categories and meanings. She identified the so-called 'paradox of discontent': if discontent is to be taken seriously societally and/or scientifically, it needs to be voiced authoritatively. However, this leads the discontented authors back to academic and therapeutic discourse, which bestows on them the authority they claim to fight.

Anyone interested in contributing to future discussions and workshops on the topic of 'philosophical activism' is much encouraged to contact one of the organizers (Laszlo Kosolosky and Gaston Meskens) and to have a look at the website.

LASZLO KOSOLOSKY
GASTON MESKENS
TOM CLAES
Ghent University

Epistemic Justification and Reasons, 1–2 November

Right at the beginning of November, "J & R"—an international workshop on "Epistemic justification and reasons"—took place at the institute of philosophy of the University of Luxembourg. The workshop was organized by Frank Hofmann (i.e., myself), and the seven invited speakers from Germany, England, Luxembourg, Belgium, and the US presented current work in epistemology and, in particular, on justification, reasons, and knowledge.

The presentations fell into three groups. The first group was concerned with the relation between justification and knowledge. Timothy Williamson presented a norm approach that leads to the conviction of the new evil demon subject. The envatted subject is guilty of not complying with the *primary* truth-related norm to only believe what one knows (and so the BIV deserved a pat on the head—if it had one, concluded Williamson). As an envatted subject, one can only comply with some *secondary* norms (like being disposed to comply with the primary norm N), but not with the primary truth-related norm(s). This yields some excuse but no justification. Chris Kelp's knowledge-first virtue epistemology was in the same spirit. He presented a general theory of competent moves and abilities. In the epistemic case, the relevant ability is a knowledge ability, and justification arises when the situation conditions are not appropriate. Finally, Clayton Littlejohn defended factualism about reasons for which one believes, against statism. He focused on basing and the causal condition that is often believed to be necessary for acting for a reason. And he proposed to distinguish between a causal explanans and a cause, in order to allow for true propositions (facts) to fulfil some causal condition.

The second group dealt with new and old evil demon scenarios. Thomas Grundmann attempted to argue that the old evil demon in fact shows that justification requires reliability, and he wanted to show that nevertheless, the new evil demon does not show that justification is independent of reliability (i.e., that it does not establish mentalism). In this way, he tried to save the new evil demon intuition, and to save reliabilism at the same time. I myself tried to argue that the new evil demon intuition is mistaken, since it confuses justification with rationality. Justification depends on factive mental states in a way in which rationality does not. Therefore, the new evil demon subject cannot have justified (first-order) beliefs but only rational ones.

The third group was concerned with immediate justification. Jim Pryor presented a detailed and careful analysis of possible defeater situations that tend to undermine immediate justification, and he tried to save immediate justification by allowing for a mild kind of incoherence. Eva Schmidt presented a positive proposal about the idea of perceptual justification by non-conceptual content of perceptual experience. A special tie between the content of experience and the content of belief is needed in order for there to be perceptual reasons for belief, in her view.

In sum, there was a lot of discussion on the new evil demon intuition, whether to accept it or not and how to deal with the various responses theoretically. Considerations about content and defeaters showed up very often, as expected. The many faces of justification are still awaiting some unification, in my view.

Frank Hofmann
Institute of philosophy, University of Luxembourg

Inferentialism in Epistemology and Philosophy of Science, 11–13 November

The fourth Madrid workshop on *New Trends in the Philosophy of Science*, organised by Jesús Zamora-Bonilla and sponsored by the Spanish government research projects "Inferentialism as social epistemology" and "Inference, causality and science", took place at UNED, Madrid, on November 11–13. This workshop's theme was *Inferentialism in Epistemology and Philosophy of Science*.

On the first day, John Norton (University of Pittsburgh) proposed a "material" dissolution of the problem of induction. He suggested that the tower metaphor of inductive support, which relies on an unjustifiable hierarchical view, should be replaced by the arch metaphor, where propositions get support both from "below" and "above". John Cantwell (Royal Institute of Technology, Stockholm) proposed an inferentialist analysis of defeasible inference. The analysis replaces the two-layer conception, where defeasible inference is defined at an epistemic level that is supra-logical with respect to the semantic level, with a single-layer conception where all inference relations are defined in terms of acceptability conditions. Ioannis Votsis (University of Düsseldorf) proposed an inferentialist account of confirmation. He argued that the semantic relations between hypothesis and evidence are sufficient to determine their confirmation relation, and defended this view against a number of objections. Samuel Fletcher (University of California Irvine) discussed the role of stability as a constraint on inference from models: inferring a property from a model is warranted only if the property is inferable from all "sufficiently similar" models. Since similarity is defined on a topology whose choice depends on the aim of inquiry, the validity of the stability criterion is contextual. Julian Reiss (Durham University) contrasted the experimentalist paradigm of inference in econometrics, which gives priority to evidence gathered from methods that resemble randomised controlled trials, and the inferentialist paradigm, which instead recommends the collection of diverse bodies of evidence. He then defended the inferentialist approach as a modified version of the hypothetico-deductive method. Kareem Khalifa (Middlebury College) and Mark Risjord (Emory University) presented a pragmatic-inferentialist account of explanaton, and argued that the account addresses several problems concerning inference to the best explanation.

On the second day, Lorenzo Casini (Munich Center for Mathematical Philosophy) presented an inferential semantics of explanatory counterfactuals. The semantics is as powerful as Woodward's interventionist semantics in dealing with counterexamples to DN. It differs from Woodward's in that it does not rely on interventions but on suppositions in a belief-revision sense. Xavier de Donato (Universidade de Santiago de Compostela) proposed an inferential account of cognitive attitudes and scientific practice, which takes inspiration from normative inferentialism and the structuralist tradition. Marion Vorms (IHPST, Paris) proposed an agent-centered approach to the content of scientific theories. Logically equivalent formulations of classical mechanics (Newtonian, Lagrangian, Hamiltonian) are not theoretically equivalent, since alternative formulations answer different explanatory questions in virtue of different inferential paths. Jaakko Kuorikoski and Samuli Pöyhönen (University of Helsinki) presented an inferentialist account of the epistemic role of simulation models. Models are inference tools, useful in our scorekeeping practice. Simulation models are particular inference tools, which work like virtual experiments. Michael Williams (Johns Hopkins University) defended an inferentialist interpretation of the theory of knowledge according to which the theory studies the concept of knowledge. Knowledge is characterised by its expressive function, viz. the willingness to defer or the invitation to defer on matters of knowledge, and its practical significance, viz. the social function of making us accountable for our use of knowledge.

On the third day, Lilia Gurova (New Bulgarian University) proposed an account of entitling explanations, that is, explanations that rely not on any actual or possible causal mechanism for the explanandum, but on the explanans entailing the possibility of the explanandum. Jesús Zamora-Bonilla and Javier González de Prado (UNED) discussed the notion of entitlement in scientific reasoning. They criticised Brandom's characterisation of induction as "entitlement-preservation". They suggested that the importance of entitlement in science lies in that scientists must decide on which inferential norms they are entitled to use for science to be a game worth playing. Mauricio Suárez (Universidad Complutense and London Institute of Philosophy) presented an inferential account of representation. The account is deflationary rather than substantive, since it analyses the meaning of representation in terms of its use, more precisely in terms of the notions of representational force and inferential capacity of the source with respect to its target. Elena Popa (Central European University, Budapest) argued that there is a hidden tension between Woodward's claim that the interventionist account of causation is concerned with causal inference but agnostic about metaphysics, and Woodward's view that our success in causal reasoning depends on mind-independent facts about causation. José Zalabardo (University College London) reconstructed the failure of the project of logical inferentialism in Wittgenstein's *Tractatus*. Wittgenstein's ambition was to infer atomic logical structure from evidence about logical consequence relations. This depends on conditions, such as contextual uniqueness, which cannot in fact be satisfied.

Lorenzo Casini
Munich Center for Mathematical Philosophy,
LMU Munich

Calls for Papers

TRUST, ARGUMENTATION, & TECHNOLOGY: special issue of *Argument and Computation*, deadline 15 December.
BELIEF CHANGE AND ARGUMENTATION THEORY: special issue of *Annals of Mathematics and Artificial Intelligence*, deadline 15 December.
PRESUPPOSITIONS: special issue of *Topoi*, deadline 15 May 2014.
VIRTUES & ARGUMENTS: special issue of *Topoi*, deadline 1 September 2014.

WHAT'S HOT IN . . .

Uncertain Reasoning

Maurice G. Kendall (1953 "The Analysis of Economic Time-Series, Part I: Prices", *Journal of the Royal Statistical Society* 116, No. 1 pp. 11–25) is often cited as the first systematic attempt at bringing empirical data to bear on the analysis of market prices. On p. 13 he notes:

> The series looks like a "wandering" one, almost as if once a week the Demon of Chance drew a random number from a symmetrical population of fixed dispersion and added it to the current price to determine the next week's price. And this, we recall, is not the behavior in some small backwater market. The data derive from the Chicago wheat market over a period of fifty years during which at least two attempts were made to corner wheat, and one might have expected the wildest irregularities in the figures.

About two decades earlier, Holbrook Working had suggested analytically that the randomness of stock prices could be due to the summation of many independent factors. Hence, he claimed, they were approximately normally distributed. This idea was anticipated in 1900 by Louis Bachelier, and refined in 1959 by M.F.M. Osborne, who favoured the log-normal distribution (thereby effectively considering prices as the result of multiplicative, rather than additive, i.i.d. variables). In the intervening five decades, the idea that stock prices are essentially random appears to have reached considerable stability. In spite, that is, of the wild disagreement about what causes the randomness in the first place, and how markets should react to it—a disagreement which is perfectly captured by the apparently incongruous assignment of the 2013 Nobel Prize in Economics to Eugene Fama and Robert Shiller.

Aidan Lyon (2013: Why are Normal Distributions Normal? *The British Journal for the Philosophy of Science* Advance Access) throws a stone in the lake by arguing that either the normal distribution is not normal, or it is normal for abnormal reasons. The normal reason as to why the normal distribution is normal, is that one can prove that the sum of a large number of independent (and identically distributed) random variables asymptotically approximates the normal distribution. This kind of result

is very central to the theory of probability, so central that Polya tagged it the *Central Limit Theorem*. Insofar as theorems can explain anything at all, the CLT is often pointed at as an explanation as to why normal distributions approximate ever so closely the bell-shaped Gaussian curve. This goes well beyond the above mentioned case of stock prices, but I find it rather odd that among the wide-ranging set of examples, from bread loaves to construction engineering, Lyon shuns market prices. He nonetheless contends that the normal line of reasoning is wrong and purports to show why this is so. I don't find his abnormal line of argument particularly clear, hence I failed to be persuaded. Yet I think that the question that gives the paper its title should be an engaging one, both for philosophically minded mathematicians and for mathematically minded philosophers. Indeed, it is very reminiscent of the sort of thing David Lewis identifies as the philosopher's duty, namely to challenge platitudes that, after all the challenging, will survive as platitudes.

The resilient platitude about the CLT is reported by de Finetti in his two-volume *Theory of Probability* where he attributes to Poincaré the following thought (with which Lyon opens his paper):

> Everyone believes it: experimentalists believing that it is a mathematical theorem, mathematicians believing that it is an empirical fact.

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

Events

December

PRIMA: 16th International Conference on Principles and Practice of Multi-Agent Systems, Dunedin, New Zealand, 1–6 December.
AIC: International Workshop on Artificial Intelligence and Cognition, Turin, Italy, 3 December.
PT&P: Proof Theory and Philosophy, Groningen, 3–5 December.
TPNC: 2nd International Conference on the Theory and Practice of Natural Computing, Cáceres, Spain, 3–5 December.
AJCAI: 26th Australasian Joint Conference on Artificial Intelligence, Dunedin, New Zealand, 3–6 December.
PhiloSci21: Challenges and Tasks, Lisbon, Portugal, 4–6 December.
Belief & Rationality: Johannes Gutenberg-Universität Mainz, 5–6 December.
Explaining without Causes: Cologne, 6–7 December.
ICDM: International Conference on Data Mining, Dallas, Texas, 8–11 December.
LPAR: Logic for Programming, Artificial Intelligence and Reasoning, Stellenbosch, South Africa, 14–19 December.

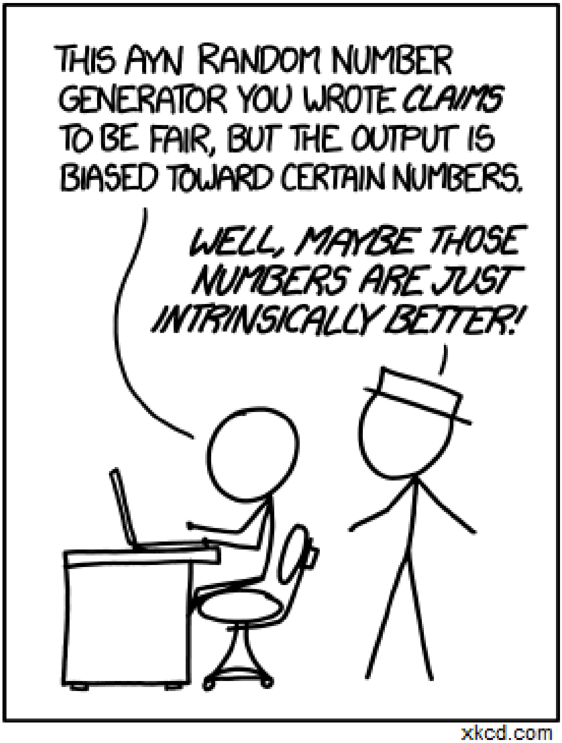

OBayes: International Workshop on Objective Bayes Methodology, Duke University, Durham, NC USA, 15–19 December.
Muskens Jubilee: Workshop in Honor of Reinhard Muskens, Tilburg University, 16 December.
Vagueness: University of Navarra, Pamplona, Spain, 16–17 December.
DialDam: 17th Workshop on the Semantics and Pragmatics of Dialogue, ILLC, University of Amsterdam, 16–18 December.
IICAI: 6th Indian International Conference on Artificial Intelligence, Tumkur, India, 18–20 December.
ISHPS: 14th Annual Conference of the Israeli Society for History & Philosophy of Science, Bloomfield Science Museum, Jerusalem, 22 December.

January

ISAIM: International Symposium on Artificial Intelligence and Mathematics, Fort Lauderdale, Florida, 6–8 January.
UUiEM: Understanding Uncertainty in Environmental Modelling, LSE, 8–10 January.

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

Conditional Thinking: Leeds, 14–15 January.
CGCPML: 7th Annual Cambridge Graduate Conference on the Philosophy of Mathematics and Logic, Cambridge, 18–19 January.

February

PHILOGICA: 3rd Colombian Conference on Logic, Epistemology, and Philosophy of Science, Bogotá, 12–14 February.
Paraconsistency: 5th World Congress on Paraconsistency, Kolkata, India, 13–17 February.
LINZ: Graded logical approaches and their applications, Linz, Austria, 18–22 February.

March

WBEM: Workshop on Beauty and Explanation in Mathematics, Umeøa University, Sweden, 11–12 March.

April

NAG: Norms, Actions, Games, London, 1–2 April.
AISB: 7th AISB Symposium on Computing and Philosophy: Is computation observer-relative?, Goldsmiths, London, 1–4 April.
HAPOP: History and Philosophy of Programming, Goldsmiths, University of London, 1–4 April.
EBL: 17th Brazilian Logic Conference, Petrópolis, Brazil, 7–11 April.
PSX4: Philosophy of Scientific Experimentation 4, Pittsburgh, PA USA, 11–12 April.
PhiloSTEM: 6th Midwest Workshop in the Philosophy of Science, Technology, Engineering, and Mathematics, Fort Wayne, Indiana, 11–12 April.
Mathematical depth: University of California, Irvine, 11–12 April.
TAMC: 11th Annual Conference on Theory and Applications of Models of Computation, Anna University, Chennai, India, 11–13 April.
PhML: Philosophy, Mathematics, Linguistics: Aspects of Interaction, St. Petersburg, Russia, 21–25 April.
PhDs in Logic: Utrecht, The Netherlands, 24–25 April.
MAICS: 25th Modern Artificial Intelligence and Cognitive Science Conference, Gonzaga University, Spokane, WA, USA, 26–27 April.
UK-CIM: UK Causal Inference Meeting (UK-CIM): Causal Inference in Health and Social Sciences, University of Cambridge, Cambridge, 28–29 April.

May

SQUARE: 4th World Congress on the Square of Opposition, Pontifical Lateran University, Vatican, 5–9 May.
MS6: Models and Simulations 6, University of Notre Dame, 9–11 May.
Formal Methods: Singapore, 12–16 May.

June

ALGMATHLOG: Algebra and Mathematical Logic: Theory and Applications, Kazan, 2–6 June.
EC: 15th ACM Conference on Economics and Computation, Stanford University, CA, USA, 8–12 June.
LOGICA: Hejnice, Czech Republic, 16–20 June.
SILFS: International Conference of the Italian Society for Logic and Philosophy of Sciences, University of Rome “Roma TRE”, 18–20 June.
AMSTA: 8th International KES Conference on Agents and Multi-agent Systems—Technologies & Applications, Crete, Greece, 18–20 June.
CiE: Computability in Europe, Budapest, Hungary, 23–27 June.

SPS: Metaphysics of Science, Lille, 25–27 June.
SPE: Semantics and Philosophy in Europe, Berlin, 26–28 June.
EGEC: 4th Annual Edinburgh Graduate Epistemology Conference, University of Edinburgh, 27–28 June.

Courses and Programmes

Courses

Modes of Technoscientific Knowledge: Chalet Giersch, Manigod, France, 19–25 January.
Groningen Winter School: Faculty of Philosophy, University of Groningen, 27–28 January.
MLSS: Machine Learning Summer School, Reykjavik, Iceland, 25 April–4 May.
Epistemic Game Theory: EPICENTER, Maastricht University, 12–23 May.
SIPTA: 6th SIPTA School on Imprecise Probabilities, Montpellier, France, 21–25 July.

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 Programmes: in Philosophy of Science, University of Leeds.
MA in Logic and Philosophy of Science: Faculty of Philosophy, Philosophy of Science and Study of Religion, LMU Munich.
MA in Logic and Theory of Science: Department of Logic of the Eotvos Lorand University, Budapest, Hungary.
MA in Metaphysics, Language, and Mind: Department of Philosophy, University of Liverpool.

MA in Mind, Brain and Learning: Westminster Institute of Education, Oxford Brookes University.
MA in Philosophy: by research, Tilburg University.
MA in Philosophy of Biological and Cognitive Sciences: Department of Philosophy, University of Bristol.
MA in Rhetoric: School of Journalism, Media and Communication, University of Central Lancashire.
MA programmes: in Philosophy of Language and Linguistics, and Philosophy of Mind and Psychology, University of Birmingham.
MRes in Cognitive Science and Humanities: Language, Communication and Organization: Institute for Logic, Cognition, Language, and Information, University of the Basque Country, Donostia, San Sebastián.
MRes in Methods and Practices of Philosophical Research: Northern Institute of Philosophy, University of Aberdeen.
MSc in Applied Statistics: Department of Economics, Mathematics and Statistics, Birkbeck, University of London.
MSc in Applied Statistics and Datamining: School of Mathematics and Statistics, University of St Andrews.
MSc in Artificial Intelligence: Faculty of Engineering, University of Leeds.

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 Sebastián).
Open Mind: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.
PhD School: in Statistics, Padua University.

Jobs and Studentships

Jobs

Associate Professor: In Philosophy of Science, University of Geneva, until filled.
Post-doc Position: in Set Theory, Torino University, until filled.
Post-doc Position: on the project "Rational reasoning with conditionals and probabilities", MCMP, LMU Munich, until filled.
Assistant Professor: in Philosophy, Department of Philosophy, Logic and Scientific Method, LSE, deadline 13 December.
Postdoc Positions: on "Science beyond scientist" project, VU Amsterdam, deadline 15 December.
Professor: of Intelligent Systems, School of Computer Science and Statistics, University College Dublin, deadline 31 January.

Studentships

Student Assistant: on the project "Rational reasoning with conditionals and probabilities", MCMP, LMU Munich, until filled.
PhD Positions: in Philosophy and Sciences of Mind, San Raffaele University, Milan, deadline 9 December.
PhD Positions: Mathematical Foundations of Computation, University of Bath, deadline 12 December.
PhD Positions: "A Study in Explanatory Power", Institute of Philosophy, University of Duisburg-Essen, deadline 13 December.
PhD Positions: Institute for Language, Cognition and Computation (ILCC), School of Informatics, University of Edinburgh, deadline 13 December.
PhD Positions: in Formal Semantics and Pragmatics, Institute for Logic, Language, and Computation (ILLC), University of Amsterdam, deadline 15 December.
PhD Position: in Logic, Rational Choice or Meta-Ethics, University of Bayreuth, deadline 15 January.