The Reasoner

Volume 5, Number 11
November 2011

www.thereasoner.org
ISSN 1757-0522

Contents

Editorial

I am happy to present this November issue of *The Reasoner*. I have chosen Wilfrid Hodges as my interview partner for this issue because of his outstanding work on formal reasoning.

In his paper "Two doors to open" (Hodges: 2006 in Gabbay, D. et al (Eds.), *Mathematical Problems from Applied Logic I: New Logics for the 21st century*, New York: Springer, pp. 277–316) Wilfrid says:

> I describe two different developments that I would like to see in logic. The first is a serious interaction between mathematical logic and cognitive science. The second is the study of the semantic ideas of medieval Arab linguists, particularly those outside the Aristotelian tradition. These two areas are very different. On the one side it seems to me that a closer cooperation between logic and cognitive science is inevitable, and the most I can hope is to nudge it along. On the other side, historical work on Arab semantics is unlikely to have any dramatic impact on present-day semantic thinking, but I'm convinced that the Arabs have things of value that should be treasured; they represent an unfamiliar viewpoint, and I hope some future workers will find it a source of inspiration.

Learning about this paper after completing the interview, I found it to be a nice coincidence that the above quote reflects the general themes of the interview: cognitive aspects of logic and the logic of Ibn Sina (Avicenna). At the beginning of the interview Wilfrid recounts how he became a model theorist. It continues with the first of the two doors to open: logic and psychology. I fully agree with Wilfrid, that researchers in both fields can learn from each other and serious interactions are highly welcome. The interview also opens the second door to the history of logic. Wilfrid tells us about thought-provoking ideas in Ibn Sina's writings, including conditionals, syllogisms and psychologism. I believe that many research areas have become highly specialized and researchers from various disciplines can learn a lot from each other. Reasoning seems to me to be a perfect area where researchers from various disciplines can meet and fruitfully interact.

Niki Pfeifer
Munich Center for Mathematical Philosophy (LMU-Munich)

Features

Interview with Wilfrid Hodges

Wilfrid Hodges is a Fellow of the British Academy and Emeritus Professor of Mathematics at Queen Mary, University of London. He has served as President of the British Logic Colloquium, ESSLLI and IUHPS/DLMPS, and is an Associate Member of the Iranian Institute of Philosophy (Tehran). He is well-known for his work in mathematical

model theory. His numerous publications deal with various topics including mathematical logic, Arabic logic, general history of logic, philosophy of mathematics, cognitive aspects of logic and music & mathematics.

Niki Pfeifer: Wilfrid, thank you for agreeing to be my interview partner for this issue of *The Reasoner*. Could you tell us something about your intellectual history? How did you become interested in logic and how did you become a model theorist?

Wilfrid Hodges: As far as I recall, I always adjusted to circumstances, and somehow circumstances landed me in model theory. The chief events as I remember them were as follows.

At school (Kings Canterbury) I studied classics as a preparation for a probable career in theology, following my parents. But I had other interests too, and in fact I spent a good deal of time reading theoretical physics. I still cherish the copy of Dirac's Quantum Physics that I bought when I was 17 or 18; I found it completely magical, and it taught me some linear algebra. I did also read some logic, though not with the same enthusiasm. When I was about 12 my father lent me a copy of Ambrose and Lazerowitz, and I remember doing the propositional logic exercises in it.

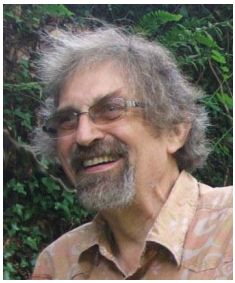

At university (New College Oxford) I followed the Hastings Rashdall course, which consisted of four years of classical languages and history and some philosophy, followed by two years of theology. I hugely enjoyed the history, which I pursued into Palestinian archaeology with Kathleen Kenyon. I'm still grateful to have read Homer and Catullus. But during that period I hardly thought at all about logic, though I do remember reading up Polish notation for an essay assignment with Tony Quinton. At the end of the six years I emerged an atheist (which was not the intention of the Hastings Rashdall course—they abolished it after me). During my last term David Wiggins, who had tutored me in philosophy, arranged for me to discuss my future with Gilbert Ryle. Ryle asked me what I most enjoyed doing, and I answered mathematics, so he said 'Call yourself a logician and do a DPhil on the back of the philosophy that you've already done, and use the course as a way get as far into mathematics as you find you can'. Later I learned that Wiggins had put him up to say this. In any case I followed Ryle's advice. In the gaps between the theology exams I bought some wire and some bazooka balls and made a free boolean algebra on two generators, to build up my boolean intuitions.

At that time John Crossley was in process of setting up a mathematical logic group in Oxford, so I joined it. The hot topic of the time was model theory, so I worked in that area. Specifically I picked up some model-theoretic ideas from Jack Silver's thesis, which C.C. Chang brought to Oxford when he visited. I worked in non-structure theory, an area that very soon afterwards became dominated by Saharon Shelah's phenomenal powers of combinatorics.

NP: Your book on mathematical logic (Chiswell, I. & Hodges, W.:2007, *Mathematical Logic*, Oxford University Press) contains two 'interludes', one on Wason's selection task and one on the Linda problem. This may seem surprising as both problems are classical reasoning tasks in psychology. I suppose these interludes are intended to point

to relations between logic and cognition. What relations do you see?

WH: Logic is, roughly speaking, the study of the formal structures involved in reasoning in a language. Since logic has been around for some time and some very clever people have contributed to it, one doesn't need to start from first principles. A good deal of logical research starts with a known formalism and develops some of its properties. Most of mathematical logic fits this description. For that kind of work, cognitive science is largely irrelevant.

But teaching is different from research. I could never teach a subject without explaining what it's for and where the principles come from. So in that book, Ian Chiswell and I give a fair number of examples of actual mathematical reasoning, and we show how to find formal patterns in them. To encourage those students who are able to think critically, we include some cases where there is a mismatch between the reasoning and the formalising of it. For example we briefly discuss donkey sentences to illustrate a use of language (even in mathematics) that doesn't formalise straightforwardly; and conversely we include the Wason and Linda examples to show that even where there is a canonical formalisation, people's thinking often doesn't follow it. We didn't have either the space or the expertise to follow up the theories that have been proposed to explain these examples. But I hope we planted a seed.

Incidentally I have no patience at all with the view of Kant, followed by Frege and some modern writers, that logic studies how we ought to think and psychology studies how we do think. A logician can tell you that if you reason by rule X, then you will sometimes find yourself deducing false conclusions from true premises. It does follow that if you want never to deduce false conclusions from true premises, you ought not to use rule X. So for example you ought not to use rule X in a research paper in pure mathematics. But in real life, where time and memory are often limited and premises are often dubious in one way or another, rule X might be for practical purposes exactly what you need. One of the major achievements of logic of the last fifty years is to start taking seriously the constraints under which we reason, and the different aims that we can have in our reasoning. This expansion of logic gives many openings for collaboration between logicians and cognitivists.

NP: This also parallels with the emergence of cognitive science in the 1950s. Cognitive science provides a platform for collaborations among various disciplines including logic and psychology. Another more recent example is the experimental philosophy movement, which might extend its current domain to reasoning research as well.

WH: There are also considerations of speed and efficiency, even in mathematical reasoning. Quine in his *Methods of Logic* describes a truth table technique which he calls 'fell swoop'; it exploits facts about truth tables that are, in Quine's words, 'visibly verifiable'. Here he is invoking a property of the human parietal lobes, though presumably he is relying on the evidence of introspection. The cognitivists have shown us that we can study the workings of our minds by much better means than pure introspection. So I hope that in the coming century sheer professionalism will lead some logicians, particularly those concerned with teaching or with practical algorithms, to pool their efforts with cognitivists.

Cognitivists and logicians belong to different communities, and there is an obvious danger that each will have a half-baked notion of what the other can offer. So we need

collaborations between the two groups. I would add that, with very few exceptions indeed, in my experience cognitivists are very willing to talk to logicians and generally have stimulating things to say.

NP: You are translating and commenting work by Ibn Sina (Avicenna). What can reasoning researchers learn from his writings?

WH: Yes, Ibn Sina from 11th century Persia. I'm going to duck this question. Making Ibn Sina's work on logic available to western researchers is going to be a major enterprise with contributions from many people. Only the first step is (nearly) complete; this was to get a reliable Arabic text. For a few decades now we have had Ibrahim Madkour's Cairo edition of the logic section of Ibn Sina's Shifa′ ('Cure'). It amounts to over two thousand pages of Arabic. There is very little point in trying to translate it until we understand better what it is about. Ibn Sina is like Carnap in that he tends to come at a question in several places from different angles, and you need to look at all the angles to get a fair picture of what he has in mind. And there are several other substantial known works of Ibn Sina in logic, not all of them published yet. Ask me again in thirty years' time—or better, ask Khaled El-Rouayheb at Harvard, who I trust will still be alive then.

I'll briefly mention some things in Ibn Sina that I have found thought-provoking. He was a competent linguist in a culture where linguistics had been developed to a deep level, and he had the great advantage of being bilingual in Arabic and Persian (from two different language families). His logical writings are dotted with remarks about linguistic usage in different languages or sciences. Several times I've asked linguists whether his observations are correct, and got a response along the lines 'That's interesting. I never thought of it but it seems to be true.' He has the same gift of observation when he discusses how people arrange their proofs, or how people are misled by metaphors.

Take for example his discussion of arguments by reductio ad absurdum. Today we carry out these arguments in a natural deduction calculus by assuming not-P, deducing a contradiction, and then 'discharging' the assumption not-P when we infer P. This device was introduced by Jaskowski and Gentzen in the 1930s. It's not clear what corresponds to 'discharging' in everyday uses of reductio ad absurdum; in fact it's not clear what constitutes making an assumption in everyday reasoning. Ibn Sina does have the notion of an assumption in the sense of a premise that we concede to somebody we are discussing something with. But he doesn't apply this notion to reductio ad absurdum. Instead he proposes that the argument is a shorthand. The propositions in the argument—or more strictly, those which in the natural deduction proof would be deduced from not-P—have an implicit 'If not-P then' in front of them. So when we deduce the contradiction, we are implicitly deducing 'If not-P then contradiction', which is equivalent to P as required. Ibn Sina notes that writers don't in fact repeat 'If not-P then' at the start of each proposition that needs it; but he regards this as a fact about normal usage. This whole account of reductio ad absurdum is remarkably close to Frege's account of conditional reasoning in some of his later writings. (But there is no line of influence from Ibn Sina to Frege; it seems none of Ibn Sina's logic was translated into any western language before the late 20th century.)

Ibn Sina sets up a Chinese wall between logic and psychology.

NP: So he holds a kind of anti-psychologism?

WH: You could say that. But he doesn't campaign against the use of psychology in logic, as Husserl and Frege did. It's more that he tries to keep the different branches of knowledge distinct in terms of the notions they use and the assumptions they make. He's particularly anxious to keep logic clear of interference from metaphysics, psychology and linguistics. So he makes remarks like 'People have a habit of spinning out the foundations of logic with things that don't belong in logic at all and are just metaphysics'. When he mentions linguistic facts, he often adds a comment to the effect 'Logicians need to know this fact, but go to a different art if you want to study the fact itself'. When he talks about logical classifications of ideas according to their types, he adds that 'investigating how an idea gets into the mind belongs to another art', i.e., psychology and not logic.

For Ibn Sina, in psychology we study how the mind works. In logic we study what the mind needs to do in order to carry out logical reasoning. Within logic we can study the algorithms, though Ibn Sina's logical writings are studiously vague about how these algorithms are implemented in the mind. His treatment of the algorithms is remarkably penetrating. For example he sets out what is almost certainly the first proof search algorithm, and very likely the first search algorithm of any kind west of India. He also notes that the main processing task in logical reasoning is to carry out a form of unification between terms, and he even has a name ('baal') for whatever part of the mind performs this task. In these cases Ibn Sina is anticipating twentieth-century ideas, but I don't know that there is much we can learn from them except some history of ideas. There is a better chance that his ideas on semantics could inspire a modern researcher, because even today we still have no definitive idea of how semantics works in natural language. I hope to have more to say about that part of Ibn Sina's work soon.

NP: Can you tell us more on Ibn Sina's work on conditionals?

WH: There is an English translation by Nabil Shehaby of the propositional logic section of Ibn Sina's Shifa′. Shehaby was a pioneer, but frankly there is a great deal in the original that we are still trying to make sense of. What follows is a very provisional sketch of how it seems to me at the moment.

The starting point for all Ibn Sina's logic is meanings of descriptive terms. I follow Jackendoff in writing for example [HORSE] for the meaning of the word 'horse'. The meaning [HORSE] is for Ibn Sina a complex object, and at its heart there is a criterion for distinguishing things that are horses from things that aren't. Sentences are like terms except that they are limited to describing situations. So [IT'S CLOUDY] contains a criterion for distinguishing situations where it's cloudy from situations where it isn't. A sentence is true if and only if there is an actual situation which fits the relevant description. (Note actual as opposed to possible. But we can't in general say 'real-world situation', because according to Ibn Sina there are mathematical truths that are not about real-world objects, unless we hold the implausible belief that somewhere in the world there is an exact regular icosahedron.)

Ibn Sina believed that we nearly always mean more than we say. For example when we say 'It's cloudy', what we mean involves a reference to the present time and place.

NP: This is also interesting in the context of conversational implicatures discussed in modern pragmatics.

WH: Yes, but I'm afraid he didn't have much to say about the mechanisms of prag-

matics. He was quite rude about some Grice-like theories of his time. His view was that even though we can and should distinguish between what a person says explicitly and what they merely imply in context, for logic this question is a distraction; a logician should work purely in terms of what the speaker meant.

It's very unclear how he thought time and place enter into the meaning of a sentence, but leave this problem aside here. The effect is that the sentence 'It's cloudy', as normally intended, is true if and only if there is a present time and place where it's cloudy. But not all sentences are intended to be just about the present; for example (one of his favourite examples—he was also a doctor) 'Going for a walk aids the digestion' is usually intended as a truth for all normal times and places. In Ibn Sina's view, most utterances involve either an implicit reference to a time and place (and in general other implied references too) or a universal or existential quantifier over times and places. So his logic, unlike Aristotle's, is basically a logic of two-quantifier sentences; he also gives examples to show that he knows of situations where more than two quantifiers are needed.

Take a standard syllogistic inference: 'Every A is a B. Every B is a C. Therefore every A is a C'. There is a procedure for drawing the conclusion from the premises; as above, it starts by unifying the B's in the two premises. Ibn Sina believed that a correct description of this procedure applies to a wider range of inferences. For example if A, B and C are sentences, we can reason 'If A then B. If B then C. Therefore if A then C'. To reach the conclusion we start by unifying the B's in the two premises, etc. etc. In Ibn Sina's terminology, the correct level for this procedure is the level of 'recombinant' inferences; in these, there are two premises that share a descriptive expression, and the conclusion is assembled from the other two descriptive expressions. This account includes both Aristotle's syllogisms and a range of conditional inferences.

NP: Mentioning Aristotle's syllogisms, is Ibn Sina making existential import assumptions? I.e., does 'Every A is a B' presuppose that A is non-empty?

WH: Yes. When there are no As, he takes affirmative sentences with subject A (for example 'Every A is a B') as false, and negative statements with subject A (for example 'It's not the case that every A is a B') as true. I don't think we know where this convention came from, but it's certainly older than Ibn Sina.

How does this convention apply to conditional statements 'If p then q'? I think on Ibn Sina's general principles he should count the 'subject' as empty if and only if p is false (or more precisely: false everywhere in the intended range of situations), and the whole sentence as affirmative if and only if q is affirmative. So suppose we look at a conditional sentence of the form 'If p then q', where q is affirmative. If the conditional is what Ibn Sina calls 'factual' (as opposed to strict implication), and is also taken to be about what is true in the present, then it should count as true whenever p is false, and thus it should be a material conditional. But there are some complications. First, even though it's very convenient to have a rule of thumb for when a sentence is true, in the last resort Ibn Sina always follows what he calls the 'mafhuum', the way a sentence is naturally understood in the standard usage of the relevant community. And second, Ibn Sina seems to think it's rather unusual in scientific reasoning (which is his main interest) to restrict oneself just to the present.

To come back to inferences: The 'recombinant' conditional inferences are pretty

much the same as the conditional inferences that Wallis in the 17th century and Boole in the 19th reduced to syllogisms. But as I read him, Ibn Sina doesn't reduce these conditional inferences to syllogisms; he thinks that they are just as basic as the syllogisms that we might reduce them to.

So how do the quantifiers over time and place appear in the conditional inferences? This is something I don't yet understand. On the face of it, when the terms are replaced by sentences, all the quantifiers should be over situations. I could say more on this, but again ask me in a few years' time. There seems to be a connection with modalities.

There is a lot more to be said about Ibn Sina's theory of conditionals, but let me mention just two points.

First, Ibn Sina believed that 'If A then B' is sometimes meant as saying that B follows from A, i.e., as a strict implication. This is a fairly standard view in the aristotelian tradition. Much less standard is Ibn Sina's claim that strict implication is not monotonic. For example if this number is five, then this number is odd. Also if this number is five and five is even, then this number is even and not odd. So far so good; but Ibn Sina also believes that adding the premise 'five is even' overrules the first inference. In other words, if this number is five and and five is even, then we can't infer that this number is odd. Classical Arabic is very bad at distinguishing between indicative and subjunctive conditionals, and I think it's clear that Ibn Sina is in fact ruling out the inference 'If this number was five and five was even, then this number would be odd'. (I slightly changed Ibn Sina's example, because his version has some unexplained features.)

Ibn Sina is well aware that this failure of monotonicity would be the kiss of death for arguments by reductio ad absurdum. He deals with this problem by concluding that in the exact sciences we use a special form of conditional which is so strong that it can't be overruled by any added premises. This seems to be a plain mistake; material implication can't be overruled by added premises, but it's the weakest form of conditional there is. So I fear the 14th century Scholastic Walter Burley was misinformed when he said that Avicenna reputedly never made a mistake in logic.

Second, Ibn Sina was aware that there are conditional inferences that don't fit the recombinant pattern. One example is modus ponens: 'A. If A then B. Therefore B.' To process this, we unify A from the two premises, but we aren't left with two other descriptive expressions to combine. Instead there is just one, which we repeat; Ibn Sina called inferences of this kind 'reduplicative'. Ibn Sina's view was that reduplicative inferences should really not be regarded as inferences. Rather they are operators which are applied to other pieces of reasoning and add a conclusion at the end. As I understand him at present, Ibn Sina hopes by this means to explain how modus ponens could have a point. Some modern writers have complained that if we already knew 'If A then B', we would have to know B in every case of A, so the inference would give us no new information. Some other writers have replied that for example 'If A then B' could be a rule that we set up by stipulation, as in 'If your dog fouls the pavement then you are liable to a fine'. Ibn Sina's view seems to be that the real work is done by the argument that the modus ponens is attached to, and the modus ponens merely rearranges the material. I suspect that his view is based at least in part on a study of concrete examples. It is a fact (which Ibn Sina duly records) that in concrete reasoning in the exact sciences, modus ponens is hardly ever spelled out.

NP: As the last question, let us turn to model theory—which will probably be one of the first associations that comes to the readers' mind if your name appears in *The Reasoner*. What do you think will be the place of model theory in the next fifty years?

WH: I'll take this to be a question about mathematical model theory, not about other related things like model-theoretic semantics or the model theory school of psychology of reasoning.

It depends on whether model theory continues to be a successful branch of mathematics. If it doesn't, it will probably carry on studying the same kinds of question as now, the same kinds of classification of first-order theories, with applications in geometry and number theory. But success in mathematics is very largely about finding things we can do that nobody ever thought of doing before. So almost by definition, we can't predict what the good mathematics of the next generation will look like. That said, the people best able to peep over the horizon are almost certainly the people who are now working close to the horizon. It's not me but my academic great-grandchildren that you should be asking.

NP: Thank you very much for this interesting interview!

WH: And many thanks to you too, Niki!

Against Cognitive Segregation: An Inseparable Connection between Creativity and Criticality

I articulate and defend the view that creative thinking and critical thinking are "so interwoven that neither can be separated from the other without an essential loss to both" (an anonymous quote cited in Richard Paul and Linda Elder: 2008, *The Thinker's Guide to the Nature and Functions of Critical and Creative Thinking*, Dillon Beach, CA: Foundation for Critical Thinking, p. 3). After sketching an influential challenge to the Inseparability Thesis, I propose a three-pronged strategy for overcoming the challenge.

Thinking, whether creative or critical, is a cognitive achievement. A deployment of mental abilities, thinking is something which the mind must work to attain. Creative thinking involves a process of molding and shaping whereas critical thinking involves a process of judging and selecting. When engaged in sound thinking, especially in a self-reflective manner, one cannot help performing the creative and critical functions simultaneously. A mature mind cannot form high-quality thoughts without ascertaining their worth, and such ascertaining includes a critical component. Nor can the mind provide well-reasoned critiques that are absent of creative effort, because it has to envision alternative possibilities and craft the analyses and interpretations involved. To produce is to appraise; to evaluate is to generate. Creativity thus requires criticality, and *vice versa*. Overlapping and interacting with each other, they are integrated and inseparable aspects of excellence of thought.

Some may oppose the Inseparability Thesis because it fails to delineate important distinctions and ultimately conflates two different ways of thinking. Opponents may invoke the following passage to justify their stance:

> Whereas creative thinking is divergent, critical thinking is convergent; whereas creative thinking tries to create something new, critical thinking

> seeks to assess worth or validity in something that exists; whereas creative thinking is carried on by violating accepted principles, critical thinking is carried on by applying accepted principles. Although creative and critical thinking may very well be different sides of the same coin, they are not identical. (Barry K. Beyer: 1987, *Practical Strategies for the Teaching of Thinking*, Allyn and Bacon, p. 35)

There are three reasons for rejecting the critique just outlined. Firstly, to maintain the Inseparability Thesis is not to identify creativity with criticality. Inseparability is not identity. That Helen of Troy is inseparable from her lover Paris does not mean that they are numerically or qualitatively identical. Pursuing the analogy one step further, we may affirm, consistently with the Inseparability Thesis, the following propositions: that while all are blessed with some measure of creativity and criticality, people tend to have more of one than the other, sometimes in the extreme; that the extent to which the former can be deliberately cultivated is more limited than the latter; and that the twin qualities can be instantiated or activated to produce entirely different and unrelated end results.

Secondly, the objection assumes that the creative/critical distinction maps neatly onto the divergent/convergent distinction, but this assumption is tendentious. Logicians, mathematicians, and scientists can engage in convergent thinking by attempting to solve problems that admit one and only one solution, and do so in a creative manner. Isaac Newton invented differential and integral calculus in order to advance his inquiries into motion and gravitation and to solve the problems thereof. Similarly, writers, composers, and artists can engage in divergent thinking by attempting to forge new ideas and new inventions, and do so in a critical manner. Leonardo da Vinci's greatness as a painter and a sculptor was due in no small part to his profound analysis and interpretation of the human anatomy. Newton's critical thinking is crucially divergent whereas da Vinci's creative thinking is crucially convergent. This conclusion is consonant with both the hypothesis that divergent thinking and convergent thinking are two major components of the creative process (J. P. Guilford: 1967, *The Nature of Human Intelligence*, McGraw-Hill) and the finding that creative people deploy both forms of thinking well. Mihaly Csikszentmihalyi, author of the finding, puts it thus:

> Divergent thinking is not much use without the ability to tell a good idea from a bad one-and this selectivity involves convergent thinking. (Csikszentmihalyi: 1996, *Creativity: Flow and the Psychology of Discovery and Invention* Harper Perennial, pp. 60–61)

Thirdly, in an essay aptly entitled "Mystery and Creation," Giorgio de Chirico, a Surrealist Italian painter, writes:

> To become truly immortal, a work of art must escape all human limits: logic and common sense will only interfere. But once these barriers are broken, it will enter the realms of childhood visions and dreams. (in *Art in Theory, 1900-2000: An Anthology of Changing Ideas*, Second Edition, ed. Charles Harrison and Paul Wood (Malden, MA: Blackwell, 2002), p. 58)

This passage is reminiscent of Beyer's above-quoted contention that

> whereas creative thinking is carried on by violating accepted principles, critical thinking is carried on by applying accepted principles.

Pace Beyer and de Chirico, even if an artist somehow succeeds in her deliberate attempt to defy logic or elude rationality in order to create a truly immortal work of art, her consciousness of what she is trying to defy or elude, together with her means-ends reasoning, indicates that whatever else it may be, her endeavor is still an exercise in critical thinking.

An inseparable connection holds between creativity and criticality. One is an indispensable component of rather than a subsidiary complement to the other. Since the acquisition, possession, retention, employment, development, and enrichment of one requires the same of the other, we cannot afford to educate our students on the unfounded assumption that creative thinking and critical thinking are segregated without compromising their ability to attain both. A cognitively integrated and holistic approach to education is recommended. Furthermore, creativity and criticality should be concurrently and conjointly cultivated across the curriculum. All else being equal, a course or program of study should not be favored or disfavored simply because it is more creative than critical, or more critical than creative, as traditionally conceived. Creative thinking and critical thinking function best in tandem.

I would like to thank Nhi Huynh, Dana McClain, Cuong Mai, Frank Williams, and three anonymous referees for their helpful comments on an earlier version of this paper.

A. Minh Nguyen
Philosophy and Religion, Eastern Kentucky University

Bolzano & Frege on Infinitely Many Truths—*a Dialogue*

In Plato's heaven Bolzano and Frege finally met; as mortals this was impossible—unfortunately. (Bolzano died in the year Frege was born: 1848.) After Gottlob Frege "walked" past the German Idealists, Bernard Bolzano came up to him. (The following is a fragment of their inspiring conversation.)

Bernard Bolzano: Hey Gottlob! Finally we meet. Aristotle introduced me to your writings here after some years. I enjoyed following your work.

Gottlob Frege: But who are you? And why are we speaking English?

BB: I am Bernard Bolzano, Husserl told me about you, but you probably did not think my work was worth reading. But now we have all the time in the world, so to speak, to discuss important matters. The English is some fashionable thing here in heaven, introduced by some philosophers of language who think that only English is philosophical—but that's dogmatic. Let us discuss truth; you made an interesting claim in your paper *der Gedanke*.

GF: You mean the claim:

> "The proposition 'I smell violets' has the same content (*Inhalt*) as the proposition 'It is true that I smell violets'. So it seems that the proposition is not altered when I add to it the property of truth" (Frege: 1918, "Der

> Gedanke," *Beitrage zur Philosophie des deutschen Idealismus* 2: p. 61—own translation)

BB: Exactly. I find this claim unintelligible. For me, if a proposition has different parts than another proposition, then the two cannot be the *same*. And even if two propositions have the same content, this does not necessarily imply that they are the same. (Bolzano, 1837: *Wissenschaftslehre* from: *Gesamtausgabe*, Stuttgart: Frommann—Holzhoog p. 448 [§96.5])

GF: Well, maybe we differ only on the definition of 'content'. Of course intentional equivalence is not sufficient for propositional identity. If so, there would be only one logical truth. But don't you agree that if 'is true' is added to a proposition, this does not alter the proposition?

BB: Certainly not! In my *Wissenschaftslehre* I wrote:

> If the proposition 'A is B' is true, then it is undeniable that the assertion 'the proposition that A is B, is true' expresses a true proposition. And this latter proposition has constituents that are different from the proposition 'A is B'; so it is a second proposition, different from the other. (Bolzano: 1837 p. 147 [§32]—own translation)

GF: But why would you claim such a thing?

BB: Well, first, I think that 'A' is the subject of 'A is B', while 'A is B' is the subject of 'A is B is true', so I want to distinguish the two; another reason is that I can come up with this argument:

1. There are only *n* truths in themselves, viz. *'A is B'*, *'C is D'*, (...) *'Y is Z'*. *[Assumption]*
2. *'A is B'*, *'C is D'*, (. . .) *'Y is Z'* are *n* truths in themselves. *[By (1)]*
3. There are no other truths, apart from *'A is B'*, *'C is D'*, (. . .) *'Y is Z'*. *[By (1)]*
4. (3) is a true proposition (or: (3) is a truth in itself). *[By (1)]*
5. There are $n + 1$ truths in themselves. *[By (2) & (4)]*
6. Steps (1)–(5) can be repeated endlessly for *n*(+1/+2/etc) truths. *[Because n is chosen arbitrarily]*
7. There are infinitely many truths in themselves. *[By (6)]*

(Based on Bolzano (1837: pp. 146–147 [§32]))

GF: Nice argument Bernhard!

BB: It's Bernard, without the "*h*". Anyway, I think you cannot use this argument, because *you* claim that the propositions 'A is B' and 'A is B is true' are the same, because they'd have the same content. The concept of 'truth' does not alter the extension for you. You thus cannot argue for the infinity of truths on the same basis.

GF: OK. Well, let me just think out loud. I could make an analogous argument on the basis of the number of propositions and then add truth to it. Thus, claim that there are only n propositions, viz. 'A is B' (...) 'Y is Z'; then I could add that this is itself a proposition. So then there are $n + 1$ propositions. I could repeat this endlessly too.

BB: Maybe, but then you'd only prove that there are infinitely many *propositions*, not *truths*. The number of true propositions may be finite, because there are infinitely many false propositions, or infinitely many propositions lacking truth-value.

GF: But since all propositions are either true or false [whether Frege actually thinks this, I am not sure, JS], the infinite many propositions would be enough: the negations of the false ones (if they'd be infinite) are true propositions. But you really think I cannot use your argument immediately?

BB: Nope. For you, 'A is B is true', is the *same* proposition as 'A is B', and thus *counting* 'A is B is true' as *another* truth next to 'A is B' is not possible. So the proposition expressed by (4) would, in your opinion, just be a *repetition* of the proposition expressed by (3), while in my argument it is, and has to be, more than that.

GF: So, because I think that adding 'is true' to a proposition does not designate another proposition, I cannot use your argument. And if I wish to argue for infinitely many truths I'd have to start with infinitely many propositions and then add 'truth' to it.

BB: Exactly. But, let us talk about your attempt to found arithmetic on logic. Were you upset when Russell discovered the paradox?

Thanks to Arianna Betti, Patricia Blanchette, Stefan Roski and two anonymous referees for helpful comments on an earlier draft of this paper. Thanks also to Marijke Langhout for making a clear representation of Bolzano's argument.

The conceptions of proposition in Bolzano and in Frege are compared and analyzed thoroughly in Künne, W., (1997) "Propositions in Bolzano and Frege," *Grazer Philosophische Studien* 53: 203–240.

Jeroen Smid
VU Amsterdam

News

Computational Logic in Multi-Agent Systems, 17–18 July

This year's edition of the workshop on *Computational Logic in Multi-Agent Systems* was held in Barcelona, Spain, alongside IJCAI. As in previous editions, CLIMA was a great place to discuss computational logic-based theories and techniques for representing, programming and reasoning about agents and multi-agent systems in a formal way.

The workshop opened with an invited talk by Simon Parsons, who proposed to use argumentation for reasoning about trust, thus connecting conclusions with individuals and sources of information. The session on *Secrets and Trust* continued with a formal analysis of information and trust propagation, and with a theoretical study of a logic of dependence between secrets.

The *Knowledge and Beliefs* session was dedicated to knowledge-based protocols, security in information exchange, modalities for modeling beliefs and information sources, belief merging, information aggregation, and the concept of definability.

Logics for Games and Social Choice were the subject of a special session, organized by Thomas Ågotnes and featuring an invited talk, in which Ulle Endriss showed many different ways in which modern logic can contribute to the study of social choice theory. Other speakers presented theoretical aspects and applications, including a new proof of the Gibbard-Satterthwaite theorem, a new logic of action and change, results on synthesising strategies for multi agent systems, a proposal for coalition logic reasoning in standard PDL-like logics, and an interpretation of ontology merging as a problem of social choice.

A session on *Cooperation* covered aspects related to interaction protocols, teams, commitment-based contracts and their exceptions, the notion of delegation, query-answering, monitoring, verification and diagnosis. Formalisms used to address these issues ranged from probabilistic model checking to logic programming, abduction and assumption-based argumentation. That was followed by three presentations on *Logic and Languages* which included a formal semantics for the Brahms modelling and simulation framework, a proposal to accommodate norm-enforcing mechanisms in the Golog agent programming language, and a proposal for rule learning, aimed to enrich abductive reasoning with a probabilistic component and to model inductive reasoning by abduction in logic programming.

The final special session on *Norms and Normative Multi-Agent Systems*, organised by Guido Boella and Leon van der Torre, started with Jan Broersen's invited talk on modeling obligations to attempt an action in a probabilistic *stit* framework extended with deontic modalities, and the effects of reasoning with probabilities on the semantics of deontic modalities. After Jan's talk, two presentations discussed actions and norms using modal style logics, and three other ones focussed on norms and normative multi-agent systems using rules style logics. These included a combination of DL ontologies with rule-based knowledge bases, a proposal to model norms as constraints to generate norm-compliant plans, and a logic of violation with time that enables the representation of reparative obligations.

The 22 papers presented at CLIMA XII and published by *Springer-Verlag* (J. Leite, P. Torroni, T. Ågotnes, G. Boella, L. van der Torre (eds.): 2011, *Computational Logic in Multi-Agent Systems, 12th International Workshop, CLIMA XII, Barcelona, Spain, July 17-18, 2011. Proceedings*, LNCS 6814. Springer-Verlag) were selected from 43 submissions by an international PC from 35 institutions located in 5 different continents. Selected and extended papers will be published in the *Journal of Logic and Computation*.

For more information about CLIMA initiatives check here.

João Leite
CENTRIA, Universidade Nova de Lisboa
Paolo Torroni
DEIS, University of Bologna

Logic, rationality and interaction, 27 August

The third LORI workshop was held in Guangzhou (China), on the campus of Sun-Yat-Sen University at the Institute of Logic and Cognition. This campus borders on the Pearl River. Already, the campus is a pleasure to be at, as it is full of large trees and beautiful shrubbery. The surrounding city blocks are a great deal less green and 20-plus mega-apartment blocks dominate them. So it is really pleasant to be at the University Campus. The riverfront is in fact an enormously long modern promenade where thousands of citizens parade, jog, and bike. This is also very pleasant.

Another distraction before talking of the content of the workshop: food. The entire LORI-3 workshop the participants were feted with banquets, for breakfast, lunch, and dinner. Every day! Now in China this is considered a normal way to treat guests, but the average European (or American) cannot but be impressed with this exceptionally hospitable treatment. What also surprised Hans and Jérôme was also that drinking red wine instead of beer has become 'de rigueur' in China in a few years. Things move quickly in China. Also, there was an excursion to the Canton Tower, a 600 meter high contraption to emit TV-signals and lure tourists. Naturally, the day of the visit the tower was more or less hidden in fog. This made it less scary to stand on the glass-bottom platform on the upper gallery at 433 meters where you can look down beyond your feet all the way to the bottom.

LORI-3 was, as the previous LORI's, a mix of invited prominent speakers, contributed accepted contributions, and accepted somewhat shorter poster contributions. The poster contributions were in a single session on Wednesday morning. About half of the poster presenters had not come to LORI—it is understandably harder to get funding if you do not have a full presentation. Maybe a reason to wonder, for the future, if this is a good way to attract participants. On the other hand, the poster presentations were well received by the audience and the presenters very well prepared for the occasion. Hans particularly liked Zuojun Xiong's able presentation about Alice forever changing her preferences about moving houses and, in the end, ending up in her original house. That seems very typical!

The invited presenters for LORI-3 were Johan van Benthem, Yongmei Liu, Ram Ramanujam, Dov Samet, Arkadii Slinko, and Kaile Su. Of course there is something nice to say about all of them. Let us simply pick out one. For Arkadii Slinko, this conference was way out of his usual research interests and for that we particularly appreciated his very successful efforts to give an overview talk on voting theory, flawlessly ending in some of his personal contributions to the area, that are about safe and unsafe strategic voting. The idea is elegant and simple. There are scenarios wherein you can call for a strategic vote and wherein you know that whatever the number of voters of your kind who follow you, the outcome cannot get worse. That is 'safe'. But there are also scenarios wherein such a block of voters can overshoot, or undershoot, its target. That is 'unsafe'. As an example, suppose A would win. But you (and many like you) prefer B. Also, you detest C. If between 10 and 50 such voters vote strategically, B will win. But if more than 50 do so, C will win, a worse outcome than A.

The main sessions of regular talks were divided by theme. There were enough talks on Dynamic Epistemic Logic to fill two sessions, no less, and there were also various

talks on the relation between knowledge and belief, and degrees of belief; and, obviously, a session on game theory related topics.

Jianying was impressed by Olivier Roy's presentation on game theory and Davide Grossi' presentation on abstract argumentation. Based on a valuable idea, which is to understand well-known game solutions not in terms of fixed informational contexts, but rather as a result of a dynamic interactive process of information exchanges, Olivier Roy presented a general characterization of interactive rationality in a game, furthermore, he gave a dynamic analysis of a well-known "paradox" arising from the admissibility algorithm in Game Theory. Davide Grossi showed that all main known semantics for abstract argumentation can be formalized in a second-order modal logic that he introduced, and presented his research on the model checking game of this logic. Moreover, he addressed the application of the game to the formalized semantics that yields adequate game-theoretic proof procedures for all known extension-based semantics, in both their skeptical and credulous versions. Taking again game theory related topics, Pavel Naumov presented some new results on the analysis of interdependencies between strategies of players in a Nash equilibrium using an independence relation between two sets of players. Based on a sound and complete axiomatization of that relation, Pavel Naumov showed that the same logical principles describe independence in four different settings: probability, information flow, concurrency and game theory.

Hans particularly recalls Yanjing Wang's in-depth presentation on axiomatic variations of public announcement logic; he was relieved that Yanjing concluded that the rule of replacement of equivalents is indeed derivable in the presentation in the textbook Dynamic Epistemic Logic, and also highly pleased to see that the composition axiom is indeed independent from the replacement of equivalents rule (and mutually so).

Jérôme enjoyed the talk by Petar Iliev on his joint work with Wiebe van der Hoek and Mike Wooldridge. The key idea is that of a system with several agents and several variables, where each agent can observe only some of the variables. A program is then executed (and this is common knowledge among the agents), which has some interesting effects on the agents' beliefs. Petar Iliav superbly explained how this framework can be useful to dining cryptographers who don't want the others to know whether or not they paid for the dinner.

We thank the ILC director Shier Ju and the ILC subdirector Minghui Xiong for their efforts and for the hospitality offered by the institute. We would also particularly like to thank Yuping Shen, who was in charge of local organization (and as such directed a whole team of helpers) and who also played an important role in the production of the proceedings

Jianying Cui
Informatics, Trier
Hans van Ditmarsch
Logic, Sevilla
Jérôme Lang
LAMSADE, Paris-Dauphin

Experimental Philosophy Group, 17–18 September

'Roll on next year's workshop—and spread the word, this is a conference not to miss!' (Paul Allonby)

The workshop was attended by academics and students from a range of disciplines, all with a shared interest in the use of experimental methods to tackle philosophical questions, or to put hypotheses based on empirical questions to the test. New insights into traditional problems and fascinating avenues for future studies are emerging from work in this developing area of philosophy.

The event started with a practical training session, led by psychologist, Kelly Schmidtke (Nottingham), on experimental design and statistical analysis, enlivened by the inclusion of experiments involving workshop participants. Pendaran Roberts (Nottingham) later presented a paper giving details of their work together on colour disagreement. The data suggest that, contrary to claims in the literature, there is no more disagreement about colour than about shape. Their findings are now being used to defend a realist view of colour against a specific objection. Paulo Sousa (QUB) gave an entertaining keynote talk presenting the results of three recent studies on the way the folk (i.e., non-philosophers) view weakness of the will, suggesting that the nature of their concept is evaluative rather than descriptive, showing parallels with ascriptions of blame or credit.

Florian Cova presented new data on the folk concept of free will from studies using subjects with frontotemporal dementia. These subjects with blunted emotions were found to be as likely to make compatibilist judgements as unimpaired subjects, challenging the view that high-affect scenarios elicit compatibilist responses because of emotional reactions.

Participants were exposed to 'bad art' during Margaret Moore's presentation. She discussed the effect of mere exposure to works of art on aesthetic appraisals, testing the hypothesis that exposure increases appreciation independently of artistic merit. The results contradicted previous findings focusing on works considered to have merit, revealing a decrease in perceived value following exposure to bad art (paintings by Kinkade). There were two discussion sessions, led by Peter Caven (Moral dilemmas and tragic remorse) and Guy Fletcher (When do we disagree?), during which participants worked together constructively to help the discussants develop proposed experimental projects, considering different methods and anticipating potential problems.

An open meeting was held, one outcome of which was the decision to develop an online forum linked to the Group's website, where researchers will be able to arrange collaborative ventures. Jonathan Webber (Cardiff) led a discussion about setting up an x-phi database.

The final talk was a thought-provoking keynote address by Joshua Knobe (Yale), discussing the folk notion of the 'true self' and exploring intuitions about cases such as when a person makes a higher-order judgement about what they want to do, but succumbs to the temptation to act otherwise. The workshop closed after a lively debate about whether the true self would generally be taken to be the one making the rational

judgement, or the one revealed in action taken in a moment of weakness.

Bryony Pierce
Department of Philosophy, University of Bristol

Causality and Explanation in the Sciences, 19–21 September

From Monday 19 till Wednesday 21 September 2011, the 'Centre for Logic and Philosophy of Science' brought philosophers and scientists together at Ghent University (Belgium) to discuss the relation between causality and explanation. This 'Causality and Explanation in the Sciences' conference (CaEitS) was already the sixth congress in the Causality in the Sciences conference series (see here).

In the first plenary session, Nancy Cartwright made the case that the effectiveness of evidence-based policy tends to suffer from locality, both with respect to its support factors (as policies do not produce results on their own) and cause-description (abstraction is needed to "get a cause that travels"). Thus construed, she argued that evidence-based policy should "mix its methods": increasing its focus on concrete details in the target and its use of cross discipline heuristics. The first day came to a close with a plenary session by Henk de Regt entitled "How we understand through causal explanation". In this presentation, it was argued that by an analysis of the concept of understanding, one can gain insight into why causal explanations provide understanding, instead of merely stipulating this.

Tuesday opened with a plenary lecture by Michael Strevens ("Causality Unified"). In his presentation, Strevens drew on his recent work on explanation to argue that the evidence adduced by the causal pluralists can be accommodated easily by a unified theory of causality, on which on all causal claims concern the same fundamental causal relation. In the plenary session that afternoon, Daniel Little argued in his talk ("Explaining the world") that social causal explanations depend upon the specification of mechanisms and processes that are at work in the social world. As there are no "laws of society" that might serve the ontological function of establishing "social necessity" for these mechanisms, Little opted to turn to features of structured human agency as the form of "necessity" that underlies causal links between antecedent conditions and the outcome in which we are interested, applying it to the example of "free-rider collective behaviour".

The final day opened with a plenary talk by Mauricio Suárez entitled "Causation, manipulability, and quantum mechanics". In his presentation, Suárez argued that, contrary to the perceived view, quantum mechanics does not constitute an exception to the applicability of the manipulability account of causal explanation, and that indeed the Causal Markov Condition is in principle applicable to the field of quantum mechanics.

Apart from the plenary sessions, around fifty contributed papers were presented. Though it is beyond the scope of this report to consider these presentations in detail, one clear division can be made, namely between those talks who focused on causality and explanation as such, and those who emphasized the application of general philosophical positions on concrete scientific disciplines. The former group consisted of contributions

on mechanistic explanations (Jon Williamson, John Pemberton, Mark Couch, Patrick McGivern, Federica Russo, Phyllis Illari, Petri Ylikoski, Ben Barros, Cyril Hédoin and Nicolas Brisset, Raoul Gervais), effect talk (Jan Willem Wieland, Alex Broadbent), causal inference (Jan Lemeire, Tim De Craecker, Frederik Van De Putte and Tjerk Gauderis, Holly Andersen, Jan Sprenger, Lorenzo Casini), interventionism (Alexandre Marcellesi, Silvia De Bianchi, Samuel Schindler), understanding (Alexandra Bradner, Wesley van Camp) and Kairetic and Structural accounts of explanation (Merel Lefevere, Alex Koo, Theo Kuipers, F.A. Muller). The latter group consisted of contributions on biology (Jan Baedke, Laszlo Kosolosky, Fridolin Gross, Michael Joffe, Leonardo Bich and Matteo Mossio), physics (Matt Farr and Alexander Reutlinger, Mark Shumelda, Michel Ghins, Andrew Wayne, Peter Bokulich), social sciences (Alessio Moneta and Tiziana Foresti, Jan Willem Lindemans, Alex Prescott-Couch, Francesca Pongiglione, Rogier De Langhe), medicine (Samantha Kleinberg, Brendan Clarke, Marshall Abrams) and mathematics (James Franklin, Victor Gijsbers, Pat Corvini, Mieke Boon).

Besides having experienced a canal boat trip through the inspiring old city centre of Ghent, the participants were treated to a number of funny clips presenting basic reasoning fallacies related to the conference topic.

Raoul Gervais
Lazslo Kosolosky
Centre for Logic and Philosophy of Science,
Ghent University

Computer Simulations and the Changing Face of Scientific Experimentation, 21–23 September

An interdisciplinary workshop on "Computer Simulations and the Changing Face of Scientific Experimentation" (see here) was held at the Simulation Technology Research Centre of the University of Stuttgart. Philosophers, historians and practising scientists from various fields came together to discuss the relation between computer simulations and experiments in contemporary science and in its recent historical development.

On the philosophical side there is consensus that simulations and experiments are distinct scientific methods. Experiments are an empirical method whereas simulations clearly are not. This restricts the possibilities for substituting experiments by simulations. For example, experiments that put fundamental hypotheses to the test can never be replaced by simulations.

While the debate about the alleged materiality of simulations that occupies a large part of the Synthese issue 2009:169 is now considered as partly misguided, questions are still open concerning the conditions under which simulations are nonetheless able to provide new knowledge about empirical systems. Anouk Barberousse (University of Lille) made a strong case for the empirical significance of simulation data, criticizing the "Laplacian view" according to which computer simulations just deterministically derive results from the premises built into them as not doing justice to the way simulation studies are conducted. In a similar vein Paul Humphreys (University of Virginia) talked about qualitative differences of data generated by simulations and experiments.

A complicated borderline case in this respect is that of empirical data that is refined by highly sophisticated computational methods as for example in CT scans.

In scientific practice simulations and experiments are regarded as distinct but complementary methods. Judith Rommel (SimTech, University of Stuttgart) in a talk on the interplay of simulations and experiments quantum chemistry and Wolfgang Nowak (SimTech), who explained how simulations are used to tune the experimental designs in hydrodynamics, demonstrated how both methods can mutually support each other if used wisely.

That empirical validation remains a crucial issue was emphasized in several talks. Sebastian Zacharias (Max-Planck-Institute Berlin) presented a scheme that classifies scientific methods according to whether they a) directly use the target object or not and b) take place in the target situation or not into four classes: lab experiments, lab simulations, field experiments, field simulations. Ultimately, what is of interest is what happens with the real target object in the field. Validation is the process of bridging the inferential gaps between simulations and experimental research or field research. This challenge was addressed by Frank Klingert and Mathias Meyer (Technical University Hamburg-Harburg) who presented a protocol for standardizing and aligning simulation research to experimental research in the field of economics. It is the field of the social sciences in particular where lack of proper validation remains a serious issue of many published simulations studies.

Several historical talks discussed the development of simulation research in the sciences in general as well as in particular branches. It seems that a periodization into three phases is now consensus among historians of science: the pre-1970 phase of slowly increasing popularity of the use of computers in science, the 1970–1990 phase of emerging high performance computing, and the post-1990 phase. Terry Shinn and Anne Markovic (University of Paris) presented empirical data that showed that the number of simulation studies virtually exploded after 1990. This data needs to be interpreted carefully, however, because in the early stages of simulation science the label "computer simulation" does not always appear in the keywords or abstracts of scientific papers even though they actually represent simulation research.

Summing it up, the workshop showed that the philosophical discussion about computer simulations has moved away from purely academic questions such as that of the "materiality of simulations" to more practice-relevant and arguably more important questions directly related to the epistemic justification of simulation research in particular contexts and for particular purposes.

Eckhart Arnold
Institute for Philosophy/SimTech,
University of Stuttgart

Social Epistemology, 25 February (Lund) and 27 September (Copenhagen)

"A formal model keeps you honest, and saves you from too much hand-waving," someone said recently. Much in this spirit, both workshops evidenced formal and informal work related to the four themes: *pluralistic ignorance*, *belief polarization*, *information cascades*, and *echo-chambers*.

Invariably, suitably pairing a formal tool and initial conditions under well-defined constraints provides ways of accounting how *prima facie* perhaps disturbing social epistemological phenomena can arise from assumptions which oftentimes seem perfectly rational. Cristina Bicchieri (UPenn) and Stephan Hartmann (Tilburg) independently demonstrated as much for the cases of *pluralistic ignorance* and *information cascades*. Varying the utilities actors assign to actions (such as 'set the trend' *vis-à-vis* 'conform to it') or Bayesian belief-updating (which decreases coherence between private beliefs and those publicly-announced) suffices to "generate" public norms, although a majority of agents does not privately endorse its content, or does so to low degree.

Jens Ulrik Hansen (Roskilde) presented various formalizations of pluralistic ignorance in an epistemic logic which is based on plausibility models. His work suggests that, from a logical point of view, pluralistic ignorance is consistent. This may go some way towards capturing that the phenomenon is robust *up until* public announcement.

As an example of information cascading, Axel Gelfert (Singapore) presented *rumor* as a species of epistemic dependence. Rumors may provide *prima facie* evidence, provided they are spread by reliable agents, who act as "filters." Their contents defeat "if so, I'd have heard" objections, whenever open flow of information cannot be assumed.

Soroush Raffiee-Rad (with Stephan Hartmann, Tilburg) modeled *anchoring effects*—e.g., the increased cognitive availability of initial and most recently received information—within Bayesian belief updating. Applied to group deliberation, agents are assumed to differ in their reliability, and able to second-order guess their interlocutors' reliability. On incremental updating, deliberative results of homogeneous groups (i.e., members are similarly reliable), e.g., expert groups, appear especially prone to anchoring.

Ryan Muldoon (UPenn) assumed each agent's perspective to provide incomplete information only (e.g., on the true state of the world). Sharing perspectives, then, improves information, which in turn might avoid or mitigate belief polarization—especially in groups that value conformity to neighboring members higher than "information obtained from the world." Based on computer simulation, such groups are shown to de-polarize, provided their learning dynamic registers a net gain through *trading perspectives*.

Sebastian Schwark (Berlin) gave hands-on examples of issues modern societies polarize over (e.g., airport runways, molecular research, carbon reduction), and outlined standard ways for political consultants to deal with them. At policy level, attitudes vary strongly with distance to one's back yard. Some groups may defy all attitude change measures (financial compensation included) *until* some policy is "put on the ground," upon which attitudes seem to change.

Presenting joint work in epistemic game theory, Eric Pacuit (Tilburg) and Olivier Roy (Munich) offered *choice rules*—e.g., avoid dominated strategies—as a normative

source for action. They outlined conditions under which the rules fail to affect players receiving "differentially good" new information, including strategic information. For instance, weakly dominated strategies may not simply count as "deleted" upon assuming it to be common knowledge that only admissible strategies are played.

Tim Kenyon (Waterloo) discussed empirical studies of groups that reach a state of *false polarization*. Here, because of biases, subgroups tend to overestimate their comparative distance apart, unless group members are instructed to consider the strengths and weaknesses of the other position (aka. 'counterfactual metacognition'). This being a promising exception, few (if any) reliable depolarization strategies are known, fewer yet remain self-administrable with the right probabilities. Studies seem to support the claim that—rather than improve *self*-discernment—teaching the bias literature (e.g., to undergraduates) merely increases the number of bias ascriptions to third parties.

Undercutting rather than tackling polarization, Mark Colyvan (Sydney) presented a *consensus guaranteed-strategy* for deliberating groups. Here, each agent assigns a weight to the priorities that govern a decision (between several goods or actions). Moreover, they weigh the degree to which they respect each other's priority-weightings. In the limit—or so a theorem states—, iterated matrix multiplication of updated priority-weights with respect-weights will terminate in consensus. Rather desirably, deliberate initial over- or underbidding is discouraged insofar as it leads to lower respect-weights.

Presentations remain available at the workshop website. Principal investigators are Vincent F. Hendricks (Copenhagen) and Erik J. Olsson (Lund). The next Copenhagen Lund Workshop in Social Epistemology takes place 9–10 December at Lund University.

Frank Zenker
Department of Philosophy, Lund University

Philosophy and Theory of Artificial Intelligence, 3–4 October

The programmatic call for the conference said: "The theory and philosophy of artificial intelligence has come to a crucial point where the agenda for the forthcoming years is in the air—this conference will try to set that agenda and to gather many of the key players." Gathering many of the key players it certainly did, with most of the 'who is who' present (1/3 of participants came to Greece from outside Europe). We had 52 speakers at the event, after a rigorous double blind reviewing procedure—see here for the program and abstracts. The conference at the leafy Anatolia campus in Thessaloniki and balmy autumn weather provided a pleasant setting with plenty of space for informal interaction, which is often the most productive part of conferences.

As for setting the agenda, at the outset of the conference we had Hubert Dreyfus warning us of the 'first step fallacy' for progress in AI (we have made the first step successfully and therefore all the next steps are straight ahead and feasible), while in the last keynote talk Jim Moor warned of a 'last step fallacy' in the area of computer ethics (we do not know whether we will ever have the last step of fully responsible artificial agents, therefore we should not make the first step towards robotic ethics). Perhaps one could use these points to characterize the agenda in the field: one camp says that

certain mental properties are necessary for intelligence and then discuss how and if these can be achieved in artificial systems, while another camp asks how we can get ahead towards more intelligent behavior without directly aiming for systems that have 'mental properties'. For the first camp, AI and Cognitive Science are just two sides of the same coin, for the second, the two are only loosely connected.

Much of the discussion at our meeting was in the first (CogSci) camp, especially on what should come after the demise of cognition as computational symbol manipulation. Aspects of non-classical Cognitive Science proposed included embodiment and morphology (Pfeifer, Gomila, Ziemke), (inter-)active cognition (Bickhard, O'Regan), dynamic systems (Dreyfus) or a new understanding of consciousness (O'Regan, Chrisley) and of meaning (Bishop, Bringsjord). (Of course, I mention only a few names—in any case, organizers see little of their own conferences.) Within that first camp, there was also a considerable current of 'classical CogSci' approaches through computational symbol manipulation or of attempts to combine both in a 'dual-theory' (e.g., Bringsjord, Bach, Gomila, Milkowski). As for the second camp of 'nonCogSci-AI', this can be seen in Pfeifer's (Brooks' style) embodiment, Sloman's 'virtual machines', Scheutz' emphasis on capabilities or in Bostrom's discussion of the consequences of upcoming machine superintelligence—and of course this camp is dominant in technical AI meetings. Finally, there are fundamental issues on the theory of computation that are central for many (C. Smith, Shagrir, Bokulich and others) and there is a move towards respectability for a discussion of 'singularity' (Dreyfus, Bostrom, Sandberg, Yampolskiy etc.). The urgency to discuss the ethics and societal relevance of AI is gaining ground—whether or not AI is ultimately 'possible'.

Overall, the theory and philosophy of AI has set itself free from the single focus on the criticism of computational symbol manipulation; it has moved towards a new Cognitive Science and, in some quarters, a less intimate link with Cognitive Science overall. These developments support a more constructive cooperation with those who do 'the real work'—but also face the real basic problems.

The conference took place at Anatolia College/ACT (Thessaloniki) and was sponsored by EUCogII, Oxford-FutureTech, AAAI, ACM-SIGART, IACAP, ECCAI. The invited papers of PT-AI 2011 will be published in two special volumes of the journal *Minds and Machines* 2012 and the section papers as a book with Springer Publishers in the new 'SAPERE' series in 2012. We have asked the participants for feedback on the event—and if that is positive (which seems likely) we will have more meetings on the "Philosophy and Theory of AI" in the future. You are welcome!
For more information see here.

Vincent C. Müller
Anatolia College/ACT &
Faculty of Philosophy, University of Oxford

Evolving Knowledge in Theory and Applications, 4 October

In recent years, intelligent agents in the contexts of open environments and multi-agent systems have become a leading paradigm in AI. Acting successfully in such environments that are uncertain, only partially accessible, and dynamic, requires sophisticated knowledge representation and reasoning techniques for the modelling of the epistemic state of the agent. In particular, in evolving environments, the agent must continuously react to new observations and to any unforeseen changes that occur. Her epistemic state must undergo corresponding changes to provide the agent with a suitable world view at any time. Thus, modern knowledge representation methods have to deal with the evolution of knowledge and belief, due to uncertain or incomplete information, or to changes in the environment.

The workshop *Evolving Knowledge in Theory and Applications* was held on October 4, 2011, in Berlin, Germany, co-located with the 34th Annual German Conference on AI (KI-2011). It was the 3rd Workshop on "Dynamics of Knowledge and Belief" (DKB-2011) organized by the Special Interest Group on Knowledge Representation and Reasoning of the Gesellschaft für Informatik (*GI-Fachgrupppe Wissensrepräsentation und Schließen*).

The particular focus of the workshop was on any topics of knowledge representation and reasoning that address the epistemic modelling of agents in open environments, and in particular on processes concerning evolving knowledge and belief, both in theory and in applications. The workshop started with a session on modelling and reasoning in probabilistic approaches.

In his talk *On Prototypical Indifference and Lifted Inference in Relational Probabilistic Conditional Logic*, Matthias Thimm investigated the complexity of probabilistic reasoning in a relational setting. Based on the notion of prototypical indifference he showed that lifted inference is no longer exponential in the number of domain elements when all predicates are unary, but is still infeasible for the general case.

Markov logic is a formalism generalising both first-order logic (for finite domains) and probabilistic graphical models. In the contribution *Knowledge Engineering with Markov Logic Networks: A Review*, Dominik Jain addressed knowledge engineering aspects with Markov logic. He described the fundamental semantics of Markov logic networks, explained how simple modelling invariants can be represented, and discussed some fallacious modelling assumptions which should be taken into account when working with these networks.

The third paper of the session dealt with notions of probabilistic inconsistencies. In *Analyzing Inconsistencies in Probabilistic Conditional Knowledge Bases using Continuous Inconsistency Measures*, Matthias Thimm discussed the problem of analyzing and measuring inconsistencies in probabilistic conditional logic by investigating inconsistency measures that support the knowledge engineer in maintaining a consistent knowledge base. He developed continuous inconsistency measures assigning a numerical value to the severity of an inconsistency which can be used for strategies for restoring consistency.

In the next session, three papers investigating aspects of relational probabilistic learning were presented. In their joint contbution *Learning Scenarios under Relational*

Probabilistic Semantics and ME Reasoning, Marc Finthammer and Nico Potyka presented a learning scenario for relational learning which takes statistics on a population as well as uncertainty on individuals into account. Developed as an extension of a propositional maximum entropy framework, it was illustrated by various examples and compared to some popular statistical relational learning approaches like Markov Logic Networks and Bayesian Logic Programs.

In *Statistical Relational Learning in Dynamic Environments—An Agent-Based Approach to Traffic Navigation Using Bayesian Logic Networks*, Daan Apeldoorn used Bayesian Logic Networks in a navigation application in a dynamic environment. Conditional probabilities are learned by an agent moving through a simulated traffic environment, and logical rules are added to determine the agent's behavior. The implementation of the agent is realized with the ProbCog Toolbox, an open-source software system for statistical relational learning.

Finally, in the talk *On Efficient Algorithms for Minimal ME-Learning*, Nico Potyka studied probabilistic learning in the context of the principle of maximum entropy. This principle states that a set of probabilistic rules is best represented by the unique probability distribution satisfying all rules and possessing maximum entropy. While in previous work, an algebraic approach was used for realizing learning by inverting maximum entropy inference, here learning was addressed with an approximative generate-and-test strategy.

The workshop attracted researchers from quite different areas of knowledge representation, thereby providing the grounds for lots of interesting discussions, and showing strong support for a follow-up workshop. The workshop proceedings and more detailed information can be found at the workshop's website.

Christoph Beierle
Mathematics and Informatics, FernUniversität Hagen
Gabriele Kern-Isberner
Computer Science, Technische Universität Dortmund

Modal and Epistemic Logic, 13–14 October

On Thursday October 13th, Jens-Ulrik Hansen defended his Ph.D. thesis "A logical toolbox for modeling knowledge and information in multi-agent systems and social epistemology" (see here).

After giving a general introduction to the thesis, Hansen gave a detailed presentation of his work on hybrid public announcement logic (HPAL) and the challenges met in constructing such a logic. One problem was that an announcement $[\varphi]$ might remove the world in which a given nominal was true, which led to various intuitive and technical problems. Hansen's solution was to use partly denoting nominals: each nominal could be true in at most one world. A sound and complete axiom system for HPAL allowing for a general completeness result was presented. A second aspect from the thesis presented was work done on terminating tableaux systems of extended hybrid logics, like HPAL and many-valued hybrid logic, another topic from the thesis. In the following question session Wiebe van der Hoek, Barteld Kooi and John Gallagher questioned Hansen on a

range of different philosophical and technical topics. This resulted in, amongst others, an interesting discussion on the interpretation of nominals in hybrid epistemic logic, where the conclusion was that these mainly made sense from the modeler's perspective, and discussion of a useful notion of bi-simulation for many-valued hybrid logics.

On Friday, October 14th, a workshop on modal and epistemic logic had been arranged. Five talks were presented, the first given by Wiebe van der Hoek on local properties in modal logic (joint work with Hans van Ditmarsch and Barteld Kooi). The motivation was to be able to add epistemic axioms like "everything a knows, b knows", $(K_a\varphi \rightarrow K_b\varphi)$, without this becoming a global validity. This was obtained by adding a modal operator, stating that a's relation was a superset of b's relation at the point of evaluation. The added operator changed bisimulation and was therefore more than "syntactic sugar". For details, refer to this AAMAS 2011 paper.

In the second talk, Valentin Goranko presented on-going research on LODIA: Logic Of Dynamics of Information and Abilities—a logical formalism taking into account both the dynamics of information and dynamics of abilities of players in multi-agent games, featuring elemnts from ATEL and DEL (joint work with Peter Hawke). The framework includes both "a priori" information regarding game structure and "empirical" information gained through play, used for strategy revision. Finally, to compute the complex semantics of the empirical information relations, Goranko suggested a semantic merge with DEL, but noted that the framework was abound with both conceptual and technical difficulties. A presentation of LODIA can be found in this LOFT 2010 paper.

Following lunch, Barteld Kooi presented joint work with Bryan Renne on Arrow Update Logic: a logic for multi-agent belief change. The basic idea behind AUL is to eliminate arrows in Kripke models when updating, rather than worlds. Kooi presented the logic and examples of use, and related it to the to DEL, showing how AUL models could be transformed into equivalent DEL action models while being (sometimes) exponentially more succinct. For more on AUL, see this paper published online in *The Review of Symbolic Logic* the day before the talk.

The second-to-last talk was given by Thomas Bolander, who presented a planning framework based on DEL. Using DEL action models, the framework shed light on issues of partial observability of the environment. Using EL to model uncertainty and action models to model actions, using an agent-internal perspective it was illustrated how the framework could successfully deal with uncertainty, nondeterminism, knowledge and multiple agents. A preprint of a joint article with Mikkel Birkegaard Andersen can be found here.

Finally, Rasmus K. Rendsvig presented ongoing work with Vincent F. Hendricks on a topic from social epistemology, informational cascades. Building on plausibility models, a merge procedure was defined for agents to obtain beliefs based on a public signal of previous agents' actions, and two cases presented, showing that cascade occurrence depends on whether the agent population considers higher-order information or not. Colorful slides are available here.

Rasmus K. Rendsvig
Roskilde University

Algorithm Game Theory, 17–19 October

The Forth International Symposium on Algorithm Game Theory (SAGT 2011) was held in Amalfi on October 17–19. The focus of this event is on the study of algorithmic aspects of game theory: typical questions include how scarce computational resources affect games and how selfishness affects the quality of the outcome in a multi-player system.

26 contributed talks and 2 invited lectures have been presented in this symposium. In the first invited talk Bruno Codenotti (Consiglio Nazionale delle Ricerche, Pisa, Italy) gave an introduction about the complexity problems related to games and equilibria; the second invited talk, held by Xiaotie Deng (University of Liverpool, Liverpool, UK), focused on markets and analyzed the two main approaches for pricing goods, namely auction mechanism and competitive market equilibrium.

As the contributed talk, I liked the talk of Laurent Gourves (Université de Paris-Dauphine, Paris, France) about "The Price of Optimum in a Matching Game", a joint work with Bruno Escoffier and Jerome Monot: it is known that there are settings where a central designer that can fix the behavior of components obtains a better outcome than if the components are distributed and selfish; in this paper, authors consider one of such settings, namely matching games on graphs, and ask how many players need to be fixed in order to have that the outcome of the game played by remaining players is the optimal outcome.

Another talk that I found really interesting was the one of Elias Koutsoupias (University of Athens, Athens, Greece) about "Scheduling without payments". This work fits in with the area of mechanism without payments, that recently attracted the attention of many game theorists and computer scientists; however, Elias was the first to develop such a mechanism for scheduling, a setting known to be hard to deal with also for mechanism with payments.

Vittorio Bilò (Università del Salento, Lecce, Italy) in his talk about "Complexity of Rational and Irrational Nash Equilibria", a joint work with Marios Mavronicolas, considered the well known problem about the rationality of a Nash equilibrium: indeed, it is known that it is possible to have a Nash equilibrium whose involved probabilities are irrational even if all utilities are expressed by rational numbers. Vittorio and Marios asks if it is possible to decide in polynomial time if a game has a rational Nash equilibria or an irrational one: unfortunately, their answer is negative.

Let me say the last words about the conference venue: Amalfi is a very pleasant place, with its mixture of sea, mountains, history and arts. To attend meetings in places like this one is always a matchless experience.

Diodato Ferraioli
Computer Science, Università di Salerno

Calls for Papers

The Alan Turing Year: special issue of *Philosophia Scientiæ*, deadline 1 November.

Computational Creativity, Intelligence, and Autonomy: special issue of *Cognitive Computation*, deadline 1 November.
Between Two Images. The Manifest and the Scientific Understanding of Man, 50 Years On: special issue of *Humana.Mente*, deadline 30 November.
Psychological Models of (Ir)rationality and Decision Making: special issue of *Synthese*, deadline 1 December.
Scope of Logic Theorems: special issue of *Logica Universalis*, deadline 24 December.
Preference Learning and Ranking: special issue of *Machine Learning*, deadline 31 December.
Structure of Scientific Revolutions: 50 Years On: special issue of *Topoi*, deadline 15 January.
Imprecision in Statistical Data Analysis: special issue of *Computational Statistics & Data Analysis*, deadline 30 January.
Formal and Intentional Semantics: special issue of *The Monist*, deadline 30 April.
The Mind-Body Problem in Cognitive Neuroscience: special issue of *Philosophia Scientiæ*, deadline 1 May.
The Aim of Belief: special issue of *Teorema*, deadline 15 September.

What's Hot in ...

...Logic and Rational Interaction

A number of new publications of interest to the readers of *The Reasoner* were announced on *LORIWEB* in the last month.

Three new papers on the dynamics of belief have appeared: Richard Booth, Thomas Meyer, Ivan Varzinczak and Renata Wassermann study belief contraction in Horn Logic; Johan van Benthem and Eric Pacuit investigate dynamic logics of evidence-based beliefs; and Sebastian Enqvist models epistemic actions in interrogative belief revision.

Two articles concern research on normative multi-agent systems: Fabio Yoshimitsu Okuyama, Rafael H. Bordini and Antonio Carlos da Rocha Costa summarize their work on situated normative infrastructures for multi-agent systems; and Mehdi Dastani, John-Jules Ch. Meyer and Davide Grossi have published work on a logic for normative multi-agent programs.

Last but not least, Johan van Benthem has made a new resource on logic and games available on his website, including many papers, and a revised version of his lecture notes on the topic that have been circulating since 2000. Van Benthem is currently preparing a monograph "Logic in Games" based on the material.

LORIWEB is a platform for sharing news related to the emerging field of Logic and Rational Interaction. If you have content to share, please contact Rasmus Rendsvig, our web manager or write to the loriweb address.

Ben Rodenhäuser
Artificial Intelligence, Groningen

. . . Uncertain Reasoning

Uncertain reasoning models fall somewhere in between the stylised pure-applied spectrum. Unlike those of pure mathematics, uncertain reasoning models are primarily evaluated according to how *useful* they are in capturing realistic patterns of rational reasoning and decision. Yet unlike those of applied mathematics, uncertain reasoning models play a fundamental role in *defining* rational reasoning and decision, the real-world phenomena they are meant to model. As a consequence, it is very hard, if not impossible, to disentangle the foundations of uncertain reasoning from the applications of its models. This peculiarity, which is well apparent in the history of the concept of probability, underlay much of the development of bayesian theory. One of the finest examples of which I am aware can be found in Leonard J. Savage's address to the Western Division of the American Philosophical Association in May 1967, which he opened as follows:

> "We statisticians, with our specific concern for uncertainty, are even more liable than other practical men to encounter philosophy, whether we like it or not. For my part I like it comparatively well. Though my background makes my knowledge and idiom somewhat different from your own." Savage (1967: Difficulties in the Theory of Personal Probability, Philosophy of Science, Vol. 34, No. 4, pp. 305-310).

After almost 45 years, such an "encounter" seems to have made it back to the top of many statisticians' agenda, yet with a dramatic change in perspective. If Savage's task was that of explaining "why a handful of statisticians who have shown competence in the past are now intent on the propaganda of indefensible and pernicious doctrine"—Bayesianism—a rapidly increasing number of practical men (and women) of our time are devoting their efforts to vindicating non-bayesian ways of doing statistics. If that shows how far the bayesian orthodoxy has gone in four decades, it also shows how little an orthodoxy is expected to remain unquestioned.

The recently established open access journal Rationality Markets and Morals is in the process of publishing a special topic (the 2.0 diachronic analogue of a classical "special issue") titled "Statistical Science and Philosophy of Science: Where Do (Should) They Meet in 2011 and Beyond?", guest edited by Deborah G. Mayo, Aris Spanos and Kent W. Staley. Among the currently available papers of this special topic, applied statistician Andrew Gelman's *Induction and Deduction in Bayesian Data Analysis* stands out as a thought-provoking contribution. The author raises a number of deep issues which he discusses, in my opinion, with uneven epistemological prudence. In an attempt to encourage readers of *The Reasoner* to delve into Gelman's own "encounter" with philosophy, I will limit myself to drawing attention to two claims which are put forward in the paper, namely (i) the orthodox view of (the philosophy of) statistics is wrong and (ii) subjectivism is bad for statistics (and science).

Gelman doesn't have too hard a time defending the rather uncontroversial thesis to the effect that "the standard view of the philosophy of statistics" is wrong and that there is much that statisticians and philosophers are missing because of this. I believe the vast majority of uncertain reasoners could not agree more, though the consensus is likely to start falling apart as soon as the details of their individual discontent with the standard

view begin to unfold. According to Gelman, the standard view is the one opposing "frequentist" and "bayesian" statistics, where the two camps are characterised in terms of a small number of key-concepts:

Frequentism	**Bayesianism**
Objective	Subjective
Procedures	Models
P-values	Bayes factors
Deduction	Induction
Falsification	Pr (model is true)

The *fil rouge* of Gelman's argument is that frequentism and bayesianism are not mutually exclusive. More specifically, the statistical practices (and beliefs) associated with each camp should not be thought of as inseparable bundles which force us to pay for unwanted gadgets. Gelman illustrates what one may naturally think of as a *pick and choose* attitude by recalling how objective bayesians make substantial use of frequentist methods to refine their choice of priors. Interestingly enough, the objective bayesian point of view on the limitation of the "standard picture" is thoroughly explored by Jon Williamson (forthcoming: "Why Frequentists and Bayesians Need Each Other", *Erkenntnis*, online first August 2011). Despite this major overlap, the two papers look at the problem from rather distinct perspectives with Gelman's focussing primarily on bayesian methods and Williamson's on bayesian theory. However, the almost simultaneous appearance of two independently motivated papers on essentially the same foundational issue clearly reaffirms the topicality of Savage's 1967 remarks.

Going back to Gelman's analysis of the standard view, he points out that one very negative consequence of subscribing to the bayesian bundle, is that it encourages bayesians to cultivate the bad habit of disregarding model checking. This, according to Gelman, depends crucially on the subjective view of probability which Bayesians embrace. One is easily reminded here of the classic line by F.J. Anscombe: "To anyone sympathetic with the current neo-Bernoullian neo-Bayesian Ramseyesque Finettist Savageous movement in statistics, the subject of testing goodness of fit is something of an embarrassment", F.J. Anscombe (1963, "Tests of Goodness of Fit", *Journal of the Royal Statistical Society*. Series B Vol. 25, No. 1, pp. 81-94). Gelman suggests that the subjectivity of the priors offers Bayesians the easy way out of such an embarrassment: since priors needn't be constrained by reality in any substantial way, the very idea of model checking can be altogether dismissed. Yet, in Gelman's opinion, this implies giving up the objectivity and ultimately the rationality of scientific enquiry. To counter this, which Gelman considers an epistemological predicament, he proposes to free bayesian modelling from subjectivity. Although the resulting "larger bayesian approach" would imply abandoning the bayesian golden standard of coherence, we should bite the bullet and tolerate some "philosophical inconsistencies" for, as argued in the last section of the paper, no uncertain reasoning model can be said to be immune to all forms of incoherence.

Gelman's contrast between bayesian subjectivism and the alleged objectivity of scientific knowledge is likely to revive a heated foundational debate which many thought had been long settled. A clear case in point is, again, Savage who extensively argued

that "the role of subjective probability in statistics is, in a sense, to make statistics less subjective", Savage (1962: *The Foundations of Statistical Inference*, Redwood Press London, p. 9).

I think that the revival of this sort of foundational question is a rather clear indication of the fact that uncertain reasoning is currently in an exciting state of fermentation.

Hykel Hosni
Centro di Ricerca Matematica, E. de Giorgi,
Scuola Normale Superiore, Pisa

Events

November

Philosophy of Medicine Roundtable: University of the Basque Country, Donostia-San Sebastian, Spain, 2–3 November.
Latin Meeting in Analytic Philosophy: Universidade de Lisboa, 2–4 November.
The Plurality of Numerical Methods in Computer Simulations and Their Philosophical Analysis: IHPST, University of Paris 1, 3–5 November.
Workshop on Ceteris Paribus Laws & Reasoning: Lund University, 4 November.
Ceteris Paribus Laws and Reasoning: Department of Philosophy, University of Lund, Sweden, 4–5 November.
CAS: Complex Adaptive Systems: Energy, Information, and Intelligence, Arlington, VA, 4–6 November.
Semantic Content: University of Barcelona, 4–6 November.
Biologically Inspired Cognitive Architectures: Arlington, Virginia, 5–6 November.
ICTAI: 23rd IEEE International Conference Tools with Artificial Intelligence, Boca Raton, Florida, USA, 7–9 November.
History and Philosophy of Computing: Celebrating the 75th anniversary of the famous 1936 Papers by A. Church, E.L. Post and A.M. Turing, Ghent University, Belgium, 7–10 November.
Ideas of Objectivity: Tübingen, 7–11 November.
Causal inference in epidemiology: London School of Hygiene & Tropical Medicine, 7–11 November.
SPR: ILCLI International Workshop on Semantics, Pragmatics, and Rhetoric, Donostia, 9–11 November.
M4M: 7th Methods for Modalities workshop, Osuna, Spain, 10–12 November.
Epistemic Norms: University of Sherbrooke, Quebec, Canada, 11–12 November.
Evolution and Norms: Concepts, Models, Challenges: Bucharest, Romania, 11–12 November.
Explanation, Causality, and Unification: Heinrich-Heine-University Düsseldorf, 11–12 November.
Reasoning with Cases in the Social Sciences: Center for Philosophy of Science, University of Pittsburgh, 11–12 November.

ACML: 3rd Asian Conference on Machine Learning, Taoyuan, Taiwan, 13–15 November.
Risk and Reliability Modelling of Energy Systems: Senate Suite, Durham Castle, 24 November.
ATAI: 2nd Annual International Conference on Advances Topics in Artificial Intelligence, Singapore, 24–25 November.
ICIIC: International Conference on Information and Intelligent Computing, Hong Kong, China, 25–27 November.
ICNI: International Conference on Networks and Information, Chengdu, China, 25–27 November.
MICAI: 10th Mexican International Conference on Artificial Intelligence, Puebla, Mexico, 26 November–4 December.
Wellington Workshop in Probability Theory and Mathematical Statistics: Victoria University, Wellington, 28–30 November.
The Roots of Deduction: Groningen, 29–30 November.
ICDeM: 1st International Conference on Decision Modeling, Kedah, Malaysia, 29 November–1 December.
Solomonoff Memorial Conference: Melbourne, Australia, 30 November–2 December.

December

CT&IT: International Workshop on Computation Theory and Information Technology, Macau, China, 1–2 December.
LENLS: Logic and Engineering of Natural Language Semantics, Takamatsu-shi, Kagawa-ken, Japan, 1–2 December.
Natural Roots of Human Cognition and Communication: Sensory-motor Concepts in Language and Science: University of Düsseldorf, Germany, 1–3 December.
ICCCI: International Conference on Computer and Computational Intelligence, Bangkok, Thailand, 2–4 December.
Indefinite Extensibility and Logical Paradoxes: Arché Research Centre, St Andrews, 2–4 December.
MindGrad: University of Warwick, UK, 3–4 December.
PT-AI: Philosophy and Theory of Artificial Intelligence, Thessaloniki, Anatolia College/ACT, 3–4 October.
NCMPL: International Conference on Non-classical Modal and Predicate Logics, Guangzhou (Canton), China, 5–9 December.
ACAL: 5th Australian Conference on Artificial Life, Perth, Murdoch, Australia, 6–8 December.
ICIRA: 4th International Conference on Intelligent Robotics and Applications, Aachen, Germany, 6–9 December.
MIWAI: 5th Multi-Disciplinary International Workshop on Artificial Intelligence, Hyderabad, Andhra Pradesh, India, 7–9 December.
The Collective Dimension of Science: Nancy, France, 8–10 December.
Copenhagen Lund Workshop in Social Epistemology: University of Lund, Sweden, 9 December.

ICACM: 1st International Conference on Advanced Computing Methodologies, Hyderabad, Andhra Pradesh, India, 9–10 December.
ICDM: 11th IEEE International Conference on Data Mining, Vancouver, Canada, 11–14 December.
IICAI: 5th Indian International Conference on Artificial Intelligence, Tumkur (near Bangalore), India, 14 December.
NIPS: 25th Annual Conference on Neural Information Processing Systems, Granada, Spain, 13–15 December.
AAL: Australasian Association of Logic, Wellington, New Zealand, 14–15 December.
COCONAT: TiLPS, Tilburg, The Netherlands, 15–16 December.
Statistics and Scientific Method I: The Controversy About Hypothesis Testing: Universidad Nacional de Educación a Distancia (UNED), Madrid, 15–16 December.
ALC: Asian Logic Colloquium, Wellington, New Zealand, 15–20 December.
Internalism versus Externalism: Universiteit van Amsterdam, 16–17 December.
Internalism versus Externalism: Institute for Logic, Language and Computation, Department of Philosophy, Universiteit van Amsterdam, 16–17 December.
ICISME: International Conference on Information Management and Systems Engineering, Nanjing, China, 16–18 December.
Computing & Statistics: Senate House, University of London, UK, 17–19 December.
Amsterdam Colloquium: ILLC, Department of Philosophy, University of Amsterdam, 19–21 December.
CAR: 3rd International Asia Conference on Informatics in Control, Automation and Robotics, Shenzhen, China, 24–25 December.

January

ISAIM: 12th International Symposium on Artificial Intelligence and Mathematics, Fort Lauderdale, Florida, 9–11 January.
University of Miami Graduate Student Conference in Epistemology: Miami, FL, 12–14 January.
Perspectivalism Workshop: Ghent, 19–20 January.
MathLog: 5th Annual Cambridge Graduate Conference on the Philosophy of Logic and Mathematics, Cambridge, 21–22 January.

February

Colombian Conference on Logic, Epistemology, and Philosophy of Science: Bogota, Colombia, 8–10 February.
Conference on Computer Science & Computational Mathematics: Melaka, Malaysia, 9–10 February.
Perspectives on Structuralism: Center for Advanced Studies (CAS) and Munich Center for Mathematical Philosophy (MCMP), LMU Munich, Germany, 16–18 February.
ICICA: International Conference on Information and Computer Applications, Hong Kong, 17–18 February.

ICCMS: 4th International Conference on Computer Modeling and Simulation, Hong Kong, 17–18 February.
Theoretical Computer Science: Auckland, New Zealand, 21–24 February.

March

FoIKS: 7th International Symposium on Foundations of Information and Knowledge Systems, Kiel, Germany, 5–9 March.
LATA: 6th International Conference on Language and Automata Theory and Applications, La Coruna, Spain, 5–9 March.
Nothing but the Truth: Vienna Forum for Analytic Philosophy, University of Vienna, 9–11 March.
LPAR: 18th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Merida, Venezuela, 11–15 March.
Axiomatic vs Semantic Truth: Munich, 14–16 March.
&HPS4: Integrated History and Philosophy of Science, Department of Philosophy and History of Science, University of Athens, 15–18 March.
Empirical Philosophy of Science. Qualitative Methods: Sandbjerg, Denmark, 21–23 March.
Pragmatism, Law, and Language: University of Idaho, 23–25 March.
New Science, New Risks: University of Pittsburgh, 30–31 March.

April

SBP: International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction, University of Maryland, 3–5 April.
Northwestern/Notre Dame Graduate Epistemology Conference: Northwestern University, Evanston, IL, 13–14 April.
Confronting Intractability in Statistical Inference: University of Bristol, 16–19 April.
Collective Intelligence: MIT, Cambridge, MA, 18–20 April.
Psychology, Emotion, and the Human Sciences: University of Windsor, Windsor, Ontario Canada, 20–21 April.
AISTATS: 15th International Conference on Artificial Intelligence and Statistics, La Palma, Canary Islands, 21–23 April.
The Progress of Science: Tilburg Center for Logic and Philosophy of Science, 25–27 April.
SDM: 12th SIAM International Conference on Data Mining, Anaheim, California, USA, 26–28 April.

May

SOPHA: Société de philosophie analytique, Paris, 4–6 May.
Belief Functions: Compiégne, France, 9–11 May.

Games, Game Theory and Game Semantics: 8th International Symposium of Cognition, Logic and Communication, Riga, Latvia, 18–20 May.
IPDPS: 26th IEEE International Parallel and Distributed Processing Symposium, Shanghai, China, 21–25 May.
JdS: 44th Journées de Statistique, Brussels, 21–25 May.
UR: Uncertain Reasoning, Special Track at FLAIRS-25, Marco Island, Florida, USA, 23–25 May.
The Aims of Inquiry and Cognition: Edinburgh Epistemology Research Group, University of Edinburgh, 25–26 May.
FEW: 9th Annual Formal Epistemology Workshop, Munich, 29 May–1 June.
Human Complexity: The University of North Carolina, Charlotte, 30 May–1 June.

June

Advances in Philosophical Logic: Ruhr University Bochum, 3–5 June.
FEW: Formal Epistemology Week, Konstanz, 4–6 June.
AAMAS: 11th International Conference on Autonomous Agents and Multiagent Systems, Valencia, Spain, 4–8 June.
Minds, Bodies, and Problems: Bilkent University, Ankara, 7–8 June.
Edinburgh Epistemology Graduate Conference: University of Edinburgh, 8–9 June.
DM: Discrete Mathematics, Dalhousie University, Halifax, Nova Scotia, Canada, 18–21 June.
LOGICA: Hejnice, northern Bohemia, 18–22 June.
CiE: Computability in Europe, University of Cambridge, Cambridge, 18–23 June.
Philosophical Insights: Senate House, University of London, 21–23 June.
HOPOS: Halifax, Nova Scotia, Canada, 21–24 June.
VaNiM: Values and Norms in Modeling, Eindhoven, The Netherlands, 25–27 June.
Square of Opposition: American University of Beirut, 26–29 June.

Courses and Programmes

Courses

SPR: ILCLI International Workshop on Semantics, Pragmatics, and Rhetoric, Institute for Logic, Cognition, Language, and Information, University of the Basque Country at Donostia, 9–11 November.
LI: Logic and Interactions, Winter School and Workshops, CIRM, Luminy, Marseille, France, 30 January–2 March.
ESSLLI: 24th European Summer School in Logic, Language and Information, Opole, Poland, 6–17 August.

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.
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 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

An interdisciplinary programme at the
University of Kent, Canterbury, UK.
Core modules provided by Philosophy and further modules from Psychology,
Computing, Statistics, Social Policy, Law, Biosciences and History.

MSc in Cognitive & Decision Sciences: Psychology, University College London.
MSc in Cognitive Science: University of Osnabrück, Germany.
MSc in Cognitive Psychology/Neuropsychology: School of Psychology, University of Kent.
MSc in Logic: Institute for Logic, Language and Computation, University of Amsterdam.
MSc in Mathematical Logic and the Theory of Computation: Mathematics, University of Manchester.
MSc in Mind, Language & Embodied Cognition: School of Philosophy, Psychology and Language Sciences, University of Edinburgh.
MSc in Philosophy of Science, Technology and Society: University of Twente, The Netherlands.
MRes in Cognitive Science and Humanities: Language, Communication and Organization: Institute for Logic, Cognition, Language, and Information, University of the Basque Country (Donostia San Sebastian).
Open Mind: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.
PhD School: in Statistics, Padua University.

Jobs and Studentships

Jobs

Post-doc positions: in Robot Learning and Reinforcement Learning, Intelligent Autonomous Systems Group, Darmstadt University of Technology / Technische Universitaet Darmstadt, Germany, to be filled asap.
Post-doc position: in the area of developmental robotics and robot learning, INRIA, Bordeaux, until filled.
Two Post-doc positions: in Machine Learning, in the project "Composing Learning for Artificial Cognitive Systems", INRIA Lille, until filled.
Post-doc position: in Machine Learning, University of Massachusetts, until filled.
Post-doc position: in Machine Learning, SUNY at Buffalo, until filled.
Post-doc position: in Philosophy of Mind, psychology, Neuroscience, and Computing, Department of Philosophy, University of Missouri-St. Louis, until filled.
Associate Professor: of Philosophy, AOS: Philosophy of Biology, Medicine, Neuroscience, Cognitive Science, and Practical Reasoning, University of Utah, until filled.
Assistant Professor: in Epistemology, Department of Philosophy, University of Pennsylvania, until filled, review of applications begins 1 November.
Invited Professor: for a one to three month visit, in the field of Probabilistic Graphical Models, Knowledge and Decision team (KOD) of the Nantes Atlantique Computer Science Lab (LINA), deadline 1 November.
Professor: in Philosophy of Science, AOS: Philosophy of biology and environmental sciences, Université du Québec à Montréal, Montreal, Canada, deadline 14 November.

Visiting Fellowships: Sydney Centre for the Foundations of Science, University of Sydney, deadline 14 November.
Assistant Professor: of Philosophy, AOS: Epistemology or Logic, at the United Arab Emirates University, deadline 15 November.
Assistant Professor: AOS: metaphysics, epistemology, and logic, Department of Philosophy, Union College, Schenectady, NY, deadline 15 November.
Assistant or Associate or Full Professor: AOS: Analytic Epistemology, Department of Philosophy and Classics, University of Texas at San Antonio, deadline 30 November.
Professor: of Artificial Intelligence, Department of Computer Science and Operations Research, University of Montreal, deadline 30 November.
Professor: of Logic, Department of Philosophy, Linguistics, and Theory of Science, University of Gothenburg, deadline 1 December.

3-year Post-Doc

To work on the relationship between Bayesian epistemology and inductive logic. Philosophy, University of Kent, deadline 15 December

Lecturer: in Logic and Philosophy of Science, Faculty of Philosophy, Louvain University, deadline 15 December.
Eight 3-year Research Fellowships: within the project "The Turing Centenary Research Project: Mind, Mechanism and Mathematics", John Templeton Foundation, deadline 16 December.
Full Professor: in High-Dimensional Data Analysis, Department of Statistics, University of South Carolina, deadline 31 December.
Professor and Tier I Canada Research Chair: in Epistemology and Metaphysics, Department of Philosophy, University of Alberta, deadline 15 January.

Studentships

PhD positions: in Robot Learning and Reinforcement Learning, Intelligent Autonomous Systems Group, Darmstadt University of Technology / Technische Universitaet Darmstadt, Germany, to be filled asap.
Three Doctoral Training Grants: School of Computing, Faculty of Engineering, University of Leeds, until filled.
One Doctoral Researcher position and one Student Research Assistant: to work in the intersection of philosophy, psychology and cognitive science, Munich Center for Mathematical Philosophy, LMU Munich, until filled.
PhD position: in the area of developmental robotics and robot learning, INRIA, Bordeaux, until filled.
Alan Musgrave Master's Scholarship: in Philosophy, University of Otago, New Zealand, deadline 1 November.
PhD position: in the Origins of Social Cognition, Centre for Philosophical Psychology, University of Antwerp, deadline 15 November.
Two PhD positions: in History of Philosophy, Ethics, Theoretical Philosophy, and Practical Philosophy, University of Groningen, deadline 16 November.

PhD studentship

To work on the relationship between Bayesian epistemology and inductive logic.
Philosophy, University of Kent, deadline 15 December

University of Kent CENTRE FOR REASONING

www.kent.ac.uk/reasoning