The Reasoner

Volume 1, Number 2
June 2007

www.thereasoner.org

Contents

§1 Editorial

There were a couple of fascinating conferences in the Netherlands this last month: Reasoning about Probabilities and Probabilistic Reasoning in Amsterdam, and Methodological Problems of the Social Sciences in Tilburg. The conferences saw a variety of disciplines come together to shed light on their principal themes. While it became clear that different disciplines offer often very different perspectives on these topics, it was equally clear that the key questions themselves—e.g., how to reason using probabilities, how to reason in the social sciences—transcend disciplinary boundaries. Why hold interdisciplinary conferences such as these? One goal is to bring potential collaborators together. More importantly, perhaps, these events increase awareness of the field, providing a source of ideas that might be transferred fruitfully to one's own discipline, and discouraging reinventions of the wheel. Moreover, these events invariably include papers that one would otherwise never encounter. Worthy goals for *The Reasoner* too.

These conferences are discussed in more detail in the News section. When you attend conferences and workshops do send in your thoughts regarding the highlights. And please keep the articles coming—submission instructions can be found on the home page http://www.thereasoner.org.

Jon Williamson, Editor
Philosophy, University of Kent

§2 Features

Causal models in cognition

Causal thinking pervades most aspects of life. Changing a light-bulb, deciding on the guilt of a crime suspect, or understanding a complex chemical process all draw on causal knowledge. Despite this ubiquity, the psychological study of causal thinking is still in its infancy. There has been plenty of theorizing about the role that causality plays in psychology, but a rigorous framework has been lacking. Such a framework is important for two reasons. First, to provide a normative benchmark against which to appraise people's causal reasoning. Second, to suggest the kinds of representations and computations that sound causal reasoning requires.

There have been numerous attempts to model human reasoning in general. These typically use either logical or probabilistic models, but neither approach captures the richness of causal inference. In particular, they have great difficulty in accommodating the distinction between intervention and observation, and the different patterns of inference that these license.

It is not surprising, therefore, that some psychologists have seized upon the recent framework of Causal Bayes networks (CBN). Here is a formal system that purports to capture the principles of causal reasoning, and offers various computational algorithms to implement it (Gopnik & Schulz, 2007: *Causal learning*, Oxford University Press; Glymour, 2001: *The mind's arrows*, MIT Press; Pearl, 2000: *Causality*, Cambridge University Press). The possibility of explaining human causal reasoning in one fell swoop is tempting indeed!

One key attribute of causal networks is that they are forward-looking; they combine representation and inference, and thus provide a means to anticipate the future rather than just summarize the past. They also model the flow of events in the world, rather than the flow of inference in the mind. Both of these properties mark out causal models as suitable representational tools in human cognition.

Another special feature of this framework is the provision of a 'calculus for action'. Causal models not only tell us what is likely to happen when one observes certain configurations of signs; they tell us what is likely to happen when one performs actions. Moreover, prediction of the effects of actions is sometimes possible even for actions that have never been performed before. Having a good causal model can thus help us anticipate the future without sticking our neck out too far. We don't have to learn by our actual mistakes, only our simulated mistakes.

The close tie between causal models and action is attractive to psychologists, but often a red-flag to philosophers. For psychologists it resonates with the idea that the fundamental purpose of cognition is to enable us to interact successfully with our environment. For philosophers it is troublesome because causality is not readily 'defined' in terms of manipulation or action, and such attempts can lead to overly restrictive or circular definitions. There are also serious concerns about several of the assumptions that underpin the CBN approach (e.g., the Causal Markov condition; the faithfulness assumption). (Cartwright, 2001: "What is Wrong with Bayes Nets?", *The Monist*, pp. 242-264; Williamson, 2005: *Bayesian nets and causality*, Oxford University Press).

Do these problems undermine the applicability of CBN to human cognition? Certainly it would seem undesirable to appraise human reasoning against a possibly flawed and over-restrictive normative theory. But then again the appropriate standard against which to appraise cognitive notions of space and time is not necessarily relativistic physics. What matters most is the fit between the reasoner's system of thinking and their environment. How does this system serve to achieve their goals? It's not fatal that CBN breaks down in special contexts, or makes restrictive assumptions, so long as it operates well in the user's typically encountered environment.

Moreover, some of the philosophical qualms can be side-stepped if we differentiate between 'methodological' and 'ultimate' principles. Several of the assumptions underlying the CBN framework might be objectionable as ultimate principles, but laudable as methodological principles used to guide our construction of causal models. These principles can be defeasible, context-dependent and restricted in their domain of applicability. Akin to the infamous principle of indifference—flawed when used to underpin definitions of probability, but fine as a maxim for intuitive probabilistic reasoning.

There are, however, more pressing practical problems. People have limited processing and storage capabilities, but large-scale Bayesian networks can be computationally demanding, especially when learning algorithms are involved. Indeed people struggle to compute or reason over three or more variables (Lagnado et al., 2007: "Beyond covariation", in Gopnik & Schulz eds. *Causal learning*, Oxford University Press). This barely allows them to compute the conditional independencies so central to CBN. And the vast amount of data typically required to identify a causal model is often unavailable to a human learner.

There are several responses to this. One is to invoke the unconscious processing power of the mind. After all, human vision is assumed to involve complex computations that would far exceed the capabilities of most conscious reasoners. And the same is supposed to apply in other areas of basic cognition. Our neurons are far smarter than we are. On this view, our (unconscious) cognitive system is well-equipped to perform the complex computations required for large-scale causal reasoning.

This move, however, relegates the conscious mind to the role of a bumbling manager, unaware what their workforce is up to. Although such a view is attractive to neuroscientists, it seems premature to give up on a proper understanding of conscious processing (and with it a significant part of human psychology). It also fails to explain why people in fact struggle to solve complex causal problems, and exhibit various simplifying biases.

A second response is to embrace the limited nature of people's processing and storage capabilities, but insist that people can still master the causal basics. They can entertain and reason with limited-capacity causal networks (possibly only a few variables at a time, depending on working memory capacity). These are constructed online, drawing on background assumptions and information from long-term memory. This can explain both sound reasoning in small-scale situations, and systematic biases in large-scale cases. Active causal networks serve as a bottleneck between information in the external world and long-term storage.

The idea that causal thinking involves limited-capacity causal models extends a key aim of CBNs: to provide simpler and more compact representations of the world. And what is lost in completeness is made up for in flexibility and speed. Simple is not only beautiful, it is also easier to reason with.

David Lagnado
Psychology, University College London

Probability and logic

It is sometimes thought that such a universal statement as ‘Every man who buys a donkey vaccinates it’ is equivalent to the associated conditional like ‘If there is a man who buys a donkey, he vaccinates it’. But they are significantly different. For the second sort of expression can be prefaced by various forms of adverbial modifiers, like ‘invariably’, ‘usually’, ‘rarely’, ‘never’, etc. And specifically the given universal statement is equivalent not to the bare conditional illustrated, but to the same conditional prefaced with ‘invariably’.

Such qualified statements are not available within the limits of standard predicate logic, but only in probability theory, and even then they do not involve subordinate statements expressible in the predicate calculus, since they have the general form

It is probable to degree n that an A is a B, i.e.

$Pr(B|A) = n$,

rather than any form like

It is probable to degree n that every A is a B,

or

It is probable to degree n that some A is a B.

The distinctive nature of the subordinate statements in conditional probability judgements is not always remarked, but Brian Ellis, in 1966, realised one very significant point when dealing with ‘a man of thirty is married’. He said of it (Ellis, B. D. 1966, *Basic Concepts of Measurement*, C.U.P. Cambridge: 170) that it “is neither true nor false. To say that it is true is to imply that all men of thirty are married. To say that it is false is to imply that no men of thirty are married. The only thing we can say is that it is probable that a man of thirty is married”. Ellis, as a result, looked forward to the day when the probability and predicate calculi would be unified into a single calculus, and in more recent years we have come to see that Hilbert’s Epsilon Calculus enables us to do just that. For subordinate statements of the required kind are available within the expressive limits of this calculus, and we can add probability operators to it.

Conditionals like that from which we started are ‘indefinite propositions’ in Ellis’ terms (Ellis, 1966: 168). Thus his simple case

If there is a man of thirty, he is married,

is

$$(\exists x)Tx \supset M\epsilon xTx,$$

and contains an epsilon term ‘ϵxTx’ (to be read in this case ‘that man of thirty’) which has what Ellis called ‘indefinite reference’. Such terms are not available within the expressive resources of standard predicate logic, but the epsilon calculus, which is a conservative extension of the predicate calculus, contains them, and within it the probability of such a conditional is a conditional probability. That is despite David Lewis’ so-called ‘triviality results’, which suggested that no conditional can have a probability that is a conditional probability (Lewis, D.K. 1976, ‘Probabilities of Conditionals and Conditional Probabilities’, *Philosophical Review*, 85: 297-315).

The point against Lewis comes from a wider consideration of certain cases of anaphoric reference, and their possible symbolisation using epsilon terms. See, for the full story, Slater, B.H. 2000: ‘Quantifier/Variable-Binding’ *Linguistics and Philosophy*, 23: 309-321, also Slater, B.H. 2004: ‘Ramsey’s Tests’ *Synthese* 141: 431-444. For epsilon terms primarily provide for the symbolisation of personal pronouns like the ‘he’ in ‘There was a man. He was talking’. This is not equivalent to ‘There was a man who was talking’, which contains, instead, the relative pronoun ‘who’. The latter sentence can be symbolised in the predicate calculus as ‘$(\exists x)(Mx.Tx)$’, but one needs the epsilon calculus to formalise the former pair of sentences as ‘$(\exists x)Mx.T\epsilon xMx$’. The possibility of formalising the cross-reference arises because epsilon terms are defined by means of the equivalence:

$$(\exists x)Fx \equiv F\epsilon xFx,$$

where the epsilon term’s reference is selected from amongst the Fs, if there are any, and from the universe at large if not. So ‘$(\exists x)Mx$’ is ‘$M\epsilon xMx$’, and ‘$T\epsilon xMx$’ continues mention of the same individual, even though the identity of that individual is indefinite.

Turning to the case of subjunctive conditionals, let us first write ‘it would be true, in world i, that p’ (i.e. what is often written in the semantics of Modal Logic ‘V(p, i) = 1’) as ‘Wip’. In connection with this we get:

$$(\exists i)Wip \equiv W[\epsilon iWip]p,$$

(where the brackets are inserted just for ease of reading), and the epsilon term 'ϵiWip', as a result, selects some world in which p is true (i.e. a 'p-world'), if there are any, and any world at all if there are not. So it gives a natural representation for the anaphoric phrase 'that case' which occurs in subjunctive conditionals like

> If there were chickens, in that case there would be eggs,

i.e.

$$(\exists i)\mathrm{Wic} \supset \mathrm{W}[\epsilon i\mathrm{Wic}]\mathrm{e}.$$

The cross reference to the antecedent world is supplied by the epsilon term in the consequent of the conditional, since that is also the epsilon term hidden in the antecedent. Notice that the natural grammar of the whole expression is also preserved, and there is no difficulty in still using hook for 'if' rather than some 'box-arrow' sign in the manner of Lewis. The subjunctive nature of the conditional is here put entirely into the content of the antecedent and consequent, as in natural language.

But the conditional so defined has a probability that is a conditional probability. For if there can be chickens the probability of

$$\neg(\exists i)\mathrm{Wic} \vee \mathrm{W}[\epsilon i\mathrm{Wic}]\mathrm{e},$$

is just the probability of the second disjunct, i.e. the chance that the chosen c-world should be an e-world. But this is just the conditional probability $pr(e|c)$. On the other hand if there cannot be chickens, the probability of the disjunction is 1, which we can take to be the conditional probability (by stipulation) in that case. In Stalnaker's rather similar theory involving a choice function there is an absurd world in which everything is true to handle this side of the matter, but with the probability of the disjunct being 1 in this case, the result is automatic.

Barry Hartley Slater
Philosophy, University of Western Australia

R v Adams and the conclusiveness of DNA evidence

DNA evidence is perhaps the most reliable evidence known to us. Yet, the *Adams* case (*R v Adams* [1996] 2 Cr App R 467) raises some difficult questions about the way DNA evidence is currently used in courts.

Adams was charged with the rape of M, to whom he was a complete stranger. M was walking home one night. The attacker approached her, asked her for the time, and then raped her from behind. She had therefore only few seconds to see his face. The prosecution case rested entirely on a DNA match between Adams and the semen from a vaginal swab. It was the first case where no evidence other than a DNA match was adduced by the prosecution. In contrast, the defence case included several pieces of evidence. First, Adams gave up his right to remain silent and gave testimony in which he denied the allegations. Second, he provided an alibi, his girlfriend who testified that he was with her at the time of the event. Third, the complainant failed to identify Adams in an identification parade. Fourth, and perhaps most importantly, Adams did not match the description that was given by M as a white, clean-shaven man with a local accent aged 20 to 25 (although at one point she said he was between 40 and 42). Nevertheless, Adams was convicted and sentenced to seven years imprisonment.

The case drew academic attention mainly for the defence's novel attempt to induce the jury to use Bayes' Theorem. The court decided that Bayes Theorem should not be used in courts because it 'plunged the jury into inappropriate and unnecessary realms of theory and complexity deflecting them from their proper tasks' (*Adams*, p. 482). The prospects and limitations of using Bayes' Theorem in courts is one of the most debated issues in the theory of evidence law (a good overview of this extensive legal debate can be found in Roberts and Zuckerman 2004, *Criminal Evidence*, OUP pp. 123-132). However, *Adams* raises some worries which are specific to the use of DNA evidence itself.

The current legal position towards DNA is somewhat one-sided in favour of the prosecution. The prosecution is allowed to adduce the fact that a DNA match was found, but also to specify the probability that 'the matching DNA characteristics would be found in the population at large' (*R v Doheny and Adams* [1997] 1 Cr App R 369, p. 369). The prosecution is even allowed to specify 'how many people with the matching characteristics are likely to be found in the United Kingdom or a more limited relevant sub-group, for instance, the caucasian, sexually active males in the Manchester area' (ibid). In *Adams*, the prosecution claimed that the probability of random match was 1 per 200 million.

The defence, on the other hand, is in a much worse position. The defence is allowed to criticise the prosecution's calculation and matching methods (ibid). Also, the judge is required to remind the jury that this probability is not the probability of guilt (ibid). However, the rule established in *Adams* prohibits the quantification of other pieces of evidence and the use of statistical methods to combine statistical and non-statistical evidence (*Adams*, p. 482). The jury is therefore left with impressive statistics but without any guidance on what to do with it.

This situation is worrying for various reasons. First, human beings tend to make systematic errors in the appreciation and analysis of statistical data (Kahneman, Slovic and Tversky 1982, *Judgment under Uncertainty: Heuristics and Biases*, CUP). For example, they tend to ignore very small probabilities and regard them as certain impossibilities. It is hard to believe that these tendencies could be overcome simply by requiring the judge to provide directions.

Second, and more importantly, the impressive figures involved might have lead the jury to underestimate the importance of other pieces of evidence (Laurence Tribe 1971, "Trial by Mathematics: Precision and Ritual in the Legal Process", 84 *Harvard Law Review* 1329). In that respect, *Adams* is especially illuminating. The defence highlighted some serious problems in the DNA evidence. For example, Dr Harris, the prosecution's expert who carried out the DNA test had 'drawn in with a pen one of the bands which had faded when reexamined' (*R v Adams* (No. 2) [1998] 1 Cr App R 377, p. 379). Whilst Dr Harris insisted that his practice 'was wholly professional' (ibid), the defence's expert statistician, Professor Donnelly, 'regarded this as a serious flaw and one which would affect any later calculation' (ibid). Professor Donnelly also criticised the size and representativeness of the database and the methods of calculation used by the prosecution's expert (ibid). Despite these difficulties and the fact that the defence adduced several pieces of counter-evidence, Adams was convicted by two separate juries. Perhaps the overwhelming impression created by the impressive figures played a vital role in this case, more than is currently realised?

One might wonder whether the current approach turns DNA evidence into *conclusive* evidence. If the DNA match in *Adams* was sufficient to prove his guilt beyond reasonable doubt, would it be ever possible for a defendant to escape conviction when DNA evidence is adduced? If the deficiencies in the DNA matching process in this case were insufficient to draw a reasonable doubt, what more is required? Furthermore, assuming that the DNA evidence was reliable, the defence evidence in *Adams* also included an alibi, the accused's testimony, lack of identification in a parade and discrepancies between the complainant's description of the rapist and the accused. Even if one thinks that all of those are insufficient to draw a reasonable doubt in the accused's guilt, is there any other hypothetical scenario in which the impressive figure of 1-per-200 million would be overcome? The *Adams* case may indicate that DNA evidence might have become conclusive. The one-sided presentation of impressive figures might leave the accused unable to draw reasonable doubt regardless what evidence he has against the DNA match or to his defence.

Amit Pundik
Law, University of Oxford

The Epistemic Basis of Legal Fact-finding

Williamson (2002: *Knowledge and Its Limits*, ch 11) has defended this connection between knowledge and assertion:

K: One must assert p only if one knows p.

Lackey challenges K in a forthcoming article ("Norms of Assertion", Noûs). She claims that knowledge cannot be what is required for proper assertion because 'there are cases in which a speaker asserts that p in the absence of knowing that p without being subject to criticism'. She gives this example: Martin, who is just beginning to see through the racist nature of the beliefs he was brought up with, is a juror at the trial of a black man charged with raping a white woman.

> After hearing the relatively flimsy evidence presented by the prosecution and the strong exculpatory evidence offered by the defense, Martin is able to recognize that the evidence clearly does not support the conclusion that the defendant committed the crime . . . In spite of this, however, he can't shake the feeling that the man . . . is guilty . . . Upon further reflection, Martin begins to suspect that such a feeling is grounded in the racism that he still harbors, and so he concludes that even if he can't quite come to believe that the defendant is innocent . . . , he nonetheless has an obligation to present the case to others this way . . . After leaving the courthouse, [a friend asks Martin] whether the "guy did it". Despite the fact that he does not believe, and hence does not know, that the defendant . . . is innocent, Martin asserts, "No, the guy did not rape her".

We are given that Martin recognizes 'that the evidence clearly does not support the conclusion that the defendant committed the crime'. This must mean that Martin knows that he does not know that the defendant is guilty. Hence, it would be wrong for Martin to tell his friend, 'Yes, the guy did rape her'. This conclusion is consistent with K.

But Martin tells his friend instead, 'No, the guy did not rape her'. On Lackey's analysis, this assertion violates K because Martin does not believe, and hence does not know, the proposition which he asserts. Nevertheless, Martin is not wrong to make the assertion; indeed, she thinks that he deserves praise rather than criticism because he is 'able to transcend his own racism and thereby offers an assertion that is both true and epistemically flawless'.

This analysis is problematic. We are told that the evidence of guilt is flimsy whereas the exculpatory evidence is strong. It is reasonable for Martin to believe at least that the defendant is probably innocent; if the evidence is strong enough, it is reasonable for him to believe that the defendant is innocent. If Martin believes neither one nor the other of these propositions, it is difficult to explain his sense of predicament. A reasonable explanation is that he is 'torn' between his belief that the defendant is (probably) innocent and his 'feeling' that the defendant is guilty. Crucially, Martin suspects that his 'feeling' is grounded in racism. While Martin finds his 'feeling' casting doubt on his belief, in recognizing the racist basis of his doubt, he prevents it from defeating his belief. (See Adler 2002, *Belief's Own Ethics* ch 10, on the compatibility of belief and doubt.) If this description of the situation is right, Martin doesn't lack the relevant belief after all.

If Martin not only believes but is justified in believing that the defendant (probably) did not commit the crime and if it is true that he is innocent, and assuming no Gettier complication exists, Martin knows that the defendant (probably) did not commit the crime. In these circumstances, Martin's assertion that 'the guy (probably) did not rape her' would be consistent with K. Whether Martin is justified in asserting the probabilistic or categorical proposition would depend on the strength of the evidence that he has. Where the evidence merely justifies the belief that the defendant probably did not commit the rape, Martin merely knows that probabilistic proposition and K only permits him to assert as much; Martin does not know that the defendant did not rape her and it would be wrong for Martin to assert categorically that he did not rape her.

Lackey proposes the following norm in place of K:

RTBNA: One should assert that p only if (i) it is reasonable for one to believe that p, and (ii) if one asserted that p, one would assert that p at least in part because it is reasonable for one to believe that p.

In the trial context, K is preferable to *RTBNA*. A guilty verdict arguably asserts that the defendant committed the crime with which he was charged. Imagine Martin is a judge sitting without a jury. He must acquit the defendant since, as noted above, it is clear that Martin knows that he does not know that the defendant is guilty; and Martin does not need to know that the defendant is innocent since an acquittal does not assert that the defendant is innocent.

Suppose the defendant was convicted on evidence which was such that it was reasonable for one to believe in his guilt, thus satisfying limb (i) of *RTBNA*, and suppose further that limb (ii) was also satisfied. As it turns out, the defendant is innocent. He was wrongfully convicted. It is beside the point that we cannot blame the court for having wrongfully convicted him. K easily explains why the conviction itself was wrongful. The guilty verdict violated K. The court asserted that the defendant was guilty when it did not know that he was guilty (and the court did not know that he was guilty because he was not). *RTBNA*, on the other hand, lacks the resources to explain why the conviction itself was wrongful.

Hock Lai Ho
Law, National University of Singapore

Why the substitution of co-referential expressions in a statement may result in change of truth-value (concluding part)

In Part I (*The Reasoner* 1(1)), I drew attention to the puzzle that, in certain contexts (attitude-ascribing contexts are the ones most commonly discussed in the literature) the inference from Fx and $x = y$ to Fy fails. No satisfying solution has been found to date. I believe that the puzzle can be solved by starting with another kind of propositional attitude context, one that has come to be known as Moore-paradoxical (after G.E. Moore who discussed it—though he was not the first to do so). Moore pointed out that there is a peculiar absurdity in someone saying something like 'It is not raining, but I believe that it is'. The paradox resides in the fact that both conjuncts could be true—so how come it is absurd to utter the conjunction of two truths??!!

Ludwig Wittgenstein's response was that to assert that p is to ostensibly express one's belief that p[1]. So the absurdity in issuing the first-person Moorean utterance is just the absurdity of expressing contradictory beliefs. But—and here's the connection with the substitutivity problem we were discussing previously—Wittgenstein did NOT hold that all utterances sharing the Moorean form were quasi-contradictions. He pointed out that in certain surroundings (Umgebungen) a Moorean utterance may make perfectly good sense. 'Imagine', says Wittgenstein, 'an announcer in a railway station who announces that a train is on schedule, but—perhaps groundlessly—is convinced that it won't arrive. He might announce: "Train No. ... will arrive at ... o'clock. Personally I don't believe it"[2]. (To bring this into line with the version of Moore's Paradox we are considering, let us have the announcer say 'Train No . . . will not arrive at . . . o'clock, but I believe that it will'.) The impishness of the announcement resides in the fact that the announcer begins by speaking (as he his paid to do) as the mouthpiece of the railway company, but suddenly switches to his own voice, distances himself from the official line, adopts his own perspective and succeeds (if the railway company is right about the train's punctuality) in uttering a truth. Without this switch of perspective half way through the sentence, his utterance would be a Moorean absurdity, a quasi-contradiction.

Normally, when reporting the views of someone, we obey a *perspectival principle*—to so frame the report as to capture the perspective of the person on whom we are reporting (the reportee). This is particularly clear in

[1] For a sensitive discussion of Wittgenstein's position, see John N. Williams, 'Wittgenstein, Moorean Absurdity and its Disappearance from Speech', *Synthese* 149/1 (2006): 225-254, see pp.232-3, and, for an explicit link to the substitutivity puzzle, pp.246-9.

[2] L.Wittgenstein, *Remarks on the Philosophy of Psychology* Vol. 1 (Oxford: Blackwell, 1980) §486). Also his *Philosophical Investigations* (Oxford: Blackwell, 1953), Part II, pp.190-2.

Valley-girl-speak (originating in California) where such reports take the form 'She was like: p' where the sentence substituting for 'p' contains the exact words *and is pronounced in the same voice* as that of the reportee. When reporting beliefs, we typically use those words with which the reportee has expressed (or, we guess, would express) those beliefs. I report to you: 'Little Annie believes that Santa Claus comes down chimneys'. On hearing this, you remind me that Santa Claus is the non-existent object named after Nicholas, a convicted German child molester. It is unremarkable that I should *not* feel obliged to infer 'Little Annie believes that the non-existent object named after Nicholas, a convicted German child molester, comes down chimneys'. To so report Little Annie's belief is in clear breach of the principle that one's report should capture the perspective of the reportee, and would amount to a distortion, a *falsification*. Little Annie does not possess the concepts of conviction and molestation, and is anyway too young to be a Meinongian.

The aforementioned principle of attitude-reporting is a defeasible one. A Grice-like principle that sometimes trumps it is that a speaker must so frame an utterance as to be intelligible to the hearer. Here is an example of such trumping: Anita tells me that she thinks her neighbour (whose name she does not know) is grotesquely fat. I wish (in conformity to Grice's maxim of *Manner*) to provide a brief report of this belief of Anita's to my close friend Harry. Harry, I know, knows Edna and knows her name but does not know what I know, namely that it is Edna who is Anita's neighbour. So I say to Harry 'Anita believes that Edna is grotesquely fat' and, *in this context* (in these surroundings) I have spoken the truth. Here I set aside the perspectival principle and say things not as Anita sees them but in such a way that Harry will understand me (in Gricean terms, I co-operate with my hearer). Had I uttered the same sentence to somebody who did not know Edna, my utterance, under those circumstances would have been *false*.

As in the Wittgenstein 'train announcement', what we have here is an example of where a sentence used in one context has a truth-value different from the truth value it has in another context (and here there are no indexicals, no ambiguous expressions and no change of referent from one context to the other). What a speaker means will depend on context and, in particular, on what he believes about his hearer's beliefs. It is in virtue of assuming knowledge shared with his hearer of such contextual features that a speaker is normally able to express himself *concisely*.

The inference with which we began this note is invalid. It needs to be replaced as follows: Let Fa be a true statement that a speaker makes in reporting the attitude of a reportee towards a. If that statement is so framed as to capture the perspective of the *reportee*, and the *reportee* believes a and b to be co-referential, and $a = b$, then Fb follows (i.e., the speaker is logically entitled to assert Fb in that context); if that statement is so framed as to capture the perspective of the *hearer* and the *hearer* believes a and b to be co-referential, and $a = b$, then Fb follows.

Laurence Goldstein
Philosophy, University of Kent

A counter on counterpossibles

In their recent contribution, "Why Counterpossibles are Non-Trivial," (2007, *The Reasoner*, 1(1), 5-6) Berit Brogaard and Joe Salerno argue that counterfactuals with impossible antecedents are not all vacuously true. Part of their paper involves critiquing an argument for the opposing 'vacuist' position advanced by Timothy Williamson. While I am sympathetic to Brogaard and Salerno's overall position, they seriously misconstrue Williamson's argument and as a result their criticism of it is ineffective. The issue is an interesting one and it is worthwhile, therefore, to take the time to properly reconstruct Williamson's argument.

Brogaard and Salerno on Williamson on counterpossibles

The argument that Brogaard and Salerno (henceforth, B&S) are concerned with comes from Williamson's book, *The Philosophy of Philosophy*. As they present it, the argument is as follows. Consider the counterpossible

(1) If 5 + 7 were 13, then 5 + 6 would be 12.

The intuitive urge is to say that (1) is non-vacuously true. But if this is right, then it also follows (non-vacuously) from the antecedent that 5 + 5 would be 11, and 5 + 4 would be 10, and . . . , and 5 + (-5) would be 1. So 0 would be 1. So if 5 + 7 were 13 then

(2) If the number of answers I gave (to a given question) were 0, then the number of answers I gave would be 1.

But (2) is plainly false. Hence the initial intuition that (1) is non-vacuously true should be rejected. On this interpretation, B&S are certainly right in their assertion that Williamson's argument is "not convincing". But the argument looks suspiciously fishy. Why should it matter whether (2) is 'plainly false'? All the non-vacuist is (allegedly) committed to is that (2) follows from the impossible antecedent, "If 5 + 7 were 13". But that some false claims follow from an impossible antecedent is hardly surprising! Thus this should not count against any particular view of counterpossibles.

Williamson's actual argument

It turns out that Williamson's actual argument is very different from the one 'reconstructed' by B&S. Williamson does not consider (1) or (2), and the intuition he aims to rebut concerns the falsity of certain counterpossibles, not their nonvacuous truth. Williamson begins by imagining a situation where he has answered '11' to the question 'What is 5 + 7?', but he mistakenly believes that he has answered '13'. So he asserts

> (3) If 5 + 7 were 13 I would have got that sum right.

Our intuition is that his utterance of (3) is false, whereas the utterance of

> (4) If 5 + 7 were 13 I would have got that sum wrong

is true. Williamson argues that this pair of intuitions cannot be correct:

> If 5 + 7 were 13 then 5 + 6 would be 12, and so (by another eleven steps) 0 would be 1, so if the number of right answers I gave were 0, the number of right answers I gave would be 1.

The structure of Williamson's argument requires a little untangling. The idea is that the above chain of reasoning reveals a contradiction between our intuition that (3) is false and our intuition that (4) is true. For we can recast (3) and (4) as

> (3*) If 5 + 7 were 13 then the number of right answers I gave would have been 1. (4*) If 5 + 7 were 13 then the number of right answers I gave would have been 0.

B&S miss the point of Williamson's argument. In their recast version, (2) has no particular relevance to (1), hence it is unclear why its derivability should matter one way or the other. Furthermore, they interpret the two sides as disagreeing not about the truth of (1) but about whether this truth is vacuous or non-vacuous. But what difference does the vacuousness or otherwise of (1) make to the derivability of (2)? But if (4*) is true, and Williamson's quoted claim is true, then (3*) follows. So—given the equivalence of (3) with (3*) and (4) with (4*)—we cannot hold both that (4) is true and (3) is false.

Other pitfalls for counterpossibles

Quite independently of the above discussion, B&S identify a second problem for Williamson's vacuist position—that in arguing for this position he himself makes use of some of the very counterpossibles whose vacuous truth he is committed to! B&S quote the following passage from Williamson:

> (5) "If all counterpossibles were false, A would be equivalent to $A \Box\!\rightarrow A$."

B&S point out that (5) is a counterpossible and thus, by Williamson's own lights, anything can truthfully be asserted as a consequent, including the negation of the claim that he makes. This is a fair point against Williamson and I shall not attempt to rebut it here. However it is worth noting that B&S themselves succumb to a similar pitfall. In support of their thesis that counterpossibles are non-trivial, B&S write

> (6) "If all counterpossibles were trivially true, much of philosophy would be less substantial than it is." (p. 6)

Now I take it that, according to B&S, (6) is a counterpossible. Thus in the situation where its antecedent is true all counterpossibles are trivially true, including (6) itself, as well as

> (7) If all counterpossibles were trivially true, all of philosophy would be just as substantial as it is.

B&S's problem is the mirror image of Williamson's. For Williamson, his general position rules out the significance of particular counterpossibles which he wishes to assert in defense of his general position. For B&S, the antecedents of some of the counterpossibles they wish to assert in attacking Williamson's position make those very counterpossibles vacuous. I conclude that neither the vacuist nor the non-vacuist position can avoid some problems of self-undermining circularity.

Alan Baker
Philosophy, Swarthmore

Pluralism about the justification of learning algorithms

It is tempting for philosophers once they have found a favoured explanation for adopting their pet mode of inductive inference to stick to it with religious fervour. Practitioners of machine learning, on the other hand, tend to be open to a range of justifications, considering the breadth of this range to be a sign of the robustness of their inductive method. In this note I would like to present four such justifications.

From the so-called 'No Free Lunch' theorem, we know that no guarantees can be given that a learning algorithm will do well on all data sets. However, there are several weaker forms of reassurance with which to avail ourselves. Here is a sample:

1. It is the only consistent scheme of inference in relation to your state of knowledge.
2. Under the assumption that the world provides you with an iid (identical independent distribution) sample from a fixed but unknown distribution, you can give an unbiased estimate of the generalisation error.
3. Under the assumption that the world provides you with an iid sample from a fixed but unknown distribution, you can give probabilistic bounds for the generalisation error in terms of the performance on a training set.
4. With no assumptions about the data generating process, you can bound online inaccuracy in terms of the error rate of the best of a specified set of competitors chosen with hindsight.

Let us consider these justifications in turn:

1) Bayesians commonly appeal to this type of justification through a variety of representation theorems. A lesser known form runs as follows (see Snoussi): my task is to find the best estimator out of a class of probability distributions over a space Z, based on a sample from Z. Best is meant here in the sense of a least 'distance' from the true distribution to my estimate, the distance often being the Kullback-Leibler divergence. Now, before I see any data, I ought to be able to describe my estimator. So let τ be my estimator, taking any finite set of data to a probability distribution. Then define the generalisation error:

$E(\tau) = \int_p P(p) \int_{\mathbf{z}} P(\mathbf{z}|p) D(p, \tau(\mathbf{z}))$,

where $D(p, .)$ is the cost of misspecifying the true distribution p. As you can see, this quantity is what you expect your loss to be without seeing the sample $\mathbf{z}$. What we seek is an estimator which minimises this expected error. Snoussi explains (p. 6) how, in the case of the Kullback-Leibler divergence, the best estimator is the function which sends $\mathbf{z}$ to $\int pP(p|\mathbf{z})$, the mean of the posterior distribution. The idea then is, if you do not behave in a Bayesian way here, you could not be said to have really had $P(p)$ as your prior.

2) The main example here is the widely used cross-validation (CV). For example, in LOO (leave one out)-CV, we take each data point of the training sample in turn, train our algorithm on the remainder of the points, and then compute the prediction of our trained algorithm on the missing point. The error rate for these predictions is an almost unbiased estimator of the generalisation error, a result due to Luntz and Brailovsky (See Vapnik, *Statistical Learning Theory*, Wiley 1998). Cross-validation is used to select model parameters. It is not specifically Bayesian, but it can be used this way.

3) Consider an algorithm whose task is to choose a classifier, either deterministic or probabilistic, on the basis of a training sample, without over-fitting. Many do so on the basis of minimising a specified combination of accuracy and simplicity, where the latter term refers to the complexity of the class from which the classifier is chosen. PAC (probably approximately correct) theorems allow us to make claims of the form: with 99% probability the gap between performance on the training data and on the test data is bounded by a certain quantity. In practice, PAC-Bayesian bounds have proved to be tightest. See my notes for more details.

4) Imagine we are fed data one point at a time. On the basis of what we have seen we must estimate a quantity relating to the most recently arrived data point. We are given a penalty according to our inaccuracy. If we adopt certain strategies, we can show that, whatever the data, our average penalty per data point will only exceed that of the best of a set of competitors, which are allowed to see all of the data at once, by at most a specified amount. This amount tends to zero as the sample size increases. In many cases, Bayesian strategies fall into this category. See Peter Grünwald's Tutorial Introduction to Minimum Description Length.

The pluralism of machine learning practitioners derives from the practical difficulties they face in their day-to-day encounters with complex learning tasks. Philosophers can profitably pay attention to this wide range of types of justification.

David Corfield
Biological Cybernetics, Tübingen

The wings are not on fire (How to turn contraposition upside down)

A Captain and the First Officer are whistling idly in an aeroplane cockpit. They are obviously very bored.

```
C:     You know what...
FO:    What?
(The Captain picks up a microphone.)
C:     (over intercom) “Hello, this is
       your Captain speaking.  There
       is absolutely no cause for
       concern.”
       That'll get them thinking.
(The First Officer reaches for the microphone.)

C:     No, no, no, no.  Not yet, not
       yet.  Let it sink in.  They'll
       be thinking, er, ‘What is there
       no cause for alarm about?  Are
       the wings on fire?’
       (over intercom) “The wings are
       not on fire.”
       Now they're thinking, er, 'why
       should he say that?'  So we
       say...
(The Steward enters.)
FO:    Oh, how are we doing?
```

```
S:    (looks down the aisle) They've
      stopped eating;
      Looking a bit worried...
C:    Good.
```

The ensuing chaos is worth a viewing of the infamous 'Airline Pilots Sketch' from *John Cleese on How to Irritate People.* (We heartily thank Michel Handgraaf for pointing this out to us.) Hilarious as it may be, the sketch also delivers a striking insight into the way logic can be turned feet over head in the simplest of all contexts—a conversation. Indeed, look closely at the reasoning underlying the Airline Pilots Sketch:

(1) If the wings are on fire, then there is cause for concern;

(2) The Captain says there is absolutely no cause for concern;

Therefore:

(3) The wings must be on fire.

Let C stand for 'There is cause for concern,' and W stand for 'The wings are on fire.' Endorsing argument 1–3 then amounts to accepting conclusion W from premises $W \supset C$ and $\neg C$. That is a logical heresy if there has ever been one. The logics of contraposition would require that we endorse a Modus Tollens argument and conclude that $\neg W$, i.e., the wings are *not* on fire. People are notoriously bad at Modus Tollens and one-third of individuals usually fail to reach any conclusion from premises of the form $P \supset Q$ and $\neg Q$. But they have never been known to reach the *opposite* of the valid conclusion, and derive P!

The fact is that they do so (Bonnefon & Villejoubert 2007: 'Modus Tollens, Modus Shmollens: Conversational effects on contrapositive reasoning', *Thinking & Reasoning, 13,* 207–222). They do so as soon as $\neg Q$ is an *assertion* rather than *a piece of information.* In other words, they do so as soon as $\neg Q$ was a statement made by someone to someone else, rather than a mere premise detached of any conversational context.

Imagine that Alice and Ben are listening to a record. You know (4).

(4) a. If Mick Jagger is singing, then this is a Rolling Stones record;
b. This is not a Rolling Stones record.

Would you conclude that Mick Jagger is singing on that record? Now imagine you know (5):

(5) a. If Mick Jagger is singing, then this is a Rolling Stones record;
b. *Alice tells Ben:* 'This is not a Rolling Stones record.'

Would you be more or less inclined to accept the conclusion that Mike Jagger is singing on the record based on (5)? On average, less than 20% people accept this conclusion from (4), but about 45% do so from (5). But why?

The reason is to be found in the way we expect people to use negation in everyday life. We do not expect people to go around randomly denying things for no reason. We do not expect a colleague to announce that 'my wife is not pregnant' if no one ever suspected she was; we do not expect the weather forecaster to announce that it will not be snowing tomorrow June 16th; and we certainly do not expect airplane pilots to routinely inform us that the wings are not (yet?) on fire. Negations in conversations are only used to contradict a belief that the speaker assumes the listener to hold, namely, a pragmatic presupposition.

Thus, as soon as someone tells us that $\neg Q$, we assume that there was some reason to think that Q was true—that is, that some P is true such that $P \supset Q$. This is the only way to preserve the assumption that the speaker is being cooperative and not wasting our time and attention with uninformative statements. But in so doing, we routinely turn logic upside down and do exactly the opposite to what proper contraposition would require. This automatic response and its potential for abuse is likely to prove an endless source of fun for TV comedians—but it is also a genuine challenge for psychological theories of deductive reasoning. It is one thing to admit that context can affect the inferences we draw, it is another to account for the fact that reasoners switch from one conclusion to its polar opposite simply as a function of whether the premises were verbally asserted or not.

Jean-François Bonnefon, Gaëlle Villejoubert
Cognition, Language(s) and Ergonomics
University of Toulouse

§3
NEWS

Reasoning about probabilities and probabilistic reasoning

Probabilistic methods are increasingly becoming an important tool in a variety of disciplines. These include computer science (probabilistic computation and automata, randomness), mathematics (probabilistic proofs), artificial intelligence (reasoning under uncertainty), epistemology (Bayesian epistemology) and linguistics (probabilistic grammars). Of course, from the beginning, probabilistic and statistical methods have been heavily used in game theory and decision theory. Often separate to discussions about applications of probabilistic methods is an important philosophical debate over the precise meaning of probabilistic and statistical statements. This debate often raises a number of issues crucial to understanding how to interpret results achieved using probabilistic methods.

Recently a conference was held at the Institute for Logic, Language and Computation at the University of Amsterdam which brought together researchers that use probabilistic and statistical methods in their respective fields and researchers concerned with the philosophical interpretation of probability and statistics to exchange ideas, approaches and techniques. The scientific and organizing committee included Horacio Arló Costa (Pittsburgh PA), Benedikt Löwe (Amsterdam), David Makinson (London), Eric Pacuit (Amsterdam) and Jan-Willem Romeijn (Groningen). There were 21 contributed talks and 10 invited talks. To get a sense of the diverse topics discussed at the conference a list of the invited speakers with titles of their talks is given below.

- Luc Bovens (London): *Dutch Books, Group Decision-Making, the Tragedy of the Commons and Strategic Jury Voting*
- David Corfield (Tübingen): *What's Happening in Machine Learning Today*
- Branden Fitelson (Berkeley CA): *Epistemological Critiques of "Classical" Logic: Two Case Studies*
- Maria Carla Galavotti (Bologna): *Probability: one or many?*
- Anne-Sophie Godfroy-Genin (Paris): *From the doctrine of probability to the theory of probabilities: the emergence of modern probability calculus*
- Peter Grünwald (Amsterdam): *Statistics without Stochastics*
- Joe Halpern (Ithaca NY): *Redoing the Foundations of Decision Theory*
- Barteld Kooi (Groningen): *Dynamic Update with Probabilities*
- Teddy Seidenfeld (Pittsburgh PA): *Concepts of Independence for Full Conditional Measures and Sets of Full Conditional Measures*

The conference also included a special session where members of the `progicnet` project presented their work. The goal of `progicnet` is to investigate the application of probabilistic networks to probabilistic logic. `progicnet` is an academic network consisting of Rolf Haenni Sola (Computer Science and Applied Mathematics, University of Bern), Jan-Willem Romeijn (Philosophy, University of Groningen), Gregory Wheeler (Artificial Intelligence, New University of Lisbon), and Jon Williamson (Philosophy, University of Kent). More information about `progicnet` can be found here. Additional information about the conference including the abstracts and full programme can be found here.

This conference is the sixth in the conference series called "Foundations of the Formal Sciences" (FotFS). This is a series of interdisciplinary conferences in mathematics, philosophy, computer science and linguistics. The main goal is to reestablish the traditionally strong links between these areas of research that have been lost in the past decades. FotFS started in 1999 as a small German workshop in Berlin. Its defining features were present from the very first meeting onwards: (a) a strong interdisciplinary spirit, (b) a focus on technical talks that nevertheless reach out to researchers from other communities, (c) a (non-exclusive) focus on young researchers. After its inaugural meeting in Berlin, FotFS was funded as a "PhD EuroConference" by the European Community, the DFG (Deutsche Forschungsgemeinschaft) and the BIGS (Bonn International Graduate School).

Each of the meetings has a distinctive topic specifying some part of the foundations of formal sciences to be investigated in an interdisciplinary way. FotFS II dealt with "Applications of Mathematical Logic in Philosophy and Linguistics", FotFS III with "Complexity in Mathematics and Computer Science", FotFS IV with "The History of the Concept of the Formal Sciences", and FotFS V with "Infinite Games". The seventh conference (FotFS VII) will be held in Brussels during fall of 2008 with the title "Bringing together philosophers and sociologists of science".

The post-proceedings of FotFS VI will be published in the series *Studies in Logic* by College Publications of King's College London and edited by B. Löwe, E. Pacuit and J.-W. Romeijn. The proceedings will contain a selection of papers from the invited and contributed speakers. The first four FotFS conferences have successfully published their proceedings:

1. FotFS I: Benedikt Löwe, Florian Rudolph (eds.), **Foundations of the Formal Sciences I**, Humboldt-Universität zu Berlin, May 7-9, 1999, special issue of Synthese (Volume 133, Number 1-2, October/November 2002).
2. FotFS II: Benedikt Löwe, Wolfgang Malzkorn, Thoralf Räsch (eds.), **Foundations of the Formal Sciences II, Applications of Mathematical Logic to Philosophy and Linguistics, Papers of a Conference held in Bonn**, November 10-13, 2000, Kluwer Academic Publishers, Dordrecht 2003 [Trends in Logic 17]
3. FotFS III: Benedikt Löwe, Boris Piwinger, Thoralf Räsch (eds.), **Classical and New Paradigms of Computation and their Complexity Hierarchies, Papers of the conference "Foundations**

of the Formal Sciences III" held in Vienna, September 21-24, 2001, Dordrecht 2004 [Trends in Logic 23]

4. FotFS IV: Benedikt Löwe, Volker Peckhaus, Thoralf Räsch (eds.), **Foundations of the Formal Sciences IV, The History of the Concept of the Formal Sciences**, College Publications, London 2006 [Studies in Logic 3]

Eric Pacuit
ILLC, University of Amsterdam

Workshop "Methodological Problems of the Social Sciences" Tilburg Institute for Logic and Philosophy of Science, May 7, 2007

On May 7, 2007, the Tilburg Institute for Logic and Philosophy of Science experienced its official opening with a workshop organised by Stephan Hartmann, the director of the institute, and the reporter. The aim of the workshop was to bring together philosophers of science and leading researchers in the social sciences to discuss current methodological problems in these sciences.

The eight talks touched on a wide variety of issues, but may be grouped roughly into four pairs: rationality in a multi-agent perspective, psychological models and methodology, normative and descriptive accounts of reasoning, and finally methodology and statistical inference.

Mark Colyvan (Philosophy, Sydney) and Stef Tijs (Econometrics, Tilburg) both took the much debated multi-agent perspective on rationality. Colyvan considered the Lehrer-Wagner model of consensus formation and proposed some modifications threreof which can guard against strategic tweaking of the consensus. Stef Tijs (Econometrics, Tilburg), one of the pioneers of the use of game theory in econometrics, discussed the practical side of game theoretic advice, showing that theoretic solutions are sometimes far removed from the negotations of firms.

Han van der Maas (Psychology, Amsterdam) and Herbert Hoijtink (Psychology, Utrecht) discussed modelling issues in psychology. Van der Maas discussed a model of intelligence inspired on ecology, in which the interaction between cognitive abilities rather than the existence of a central latent ability explains the positive correlation between cognitive ability indicators. Hoijtink showed how background knowledge on such abilities can inform the statistical procedures used to investigate them, by encoding this knowledge in a prior probability over statistical models.

The talks of Branden Fitelson (Philosophy, Berkeley) and Michiel van Lambalgen (Logic and Cognitive Science, Amsterdam) both directly concerned the relation between logic and the psychology of reasoning. Fitelson pointed to an old discussion in confirmation theory concerning the difference between relevance measures and absolute measures of confirmation, vindicating Popper in passing, and then showing that the very same distinction between these measures of confirmation can be used to explain the conjunction fallacy, probably known best from Tversky and Kahneman's example on Suzy the bankteller.

Michiel van Lambalgen provided a ciriticism of probabilistic models of reasoning. He focussed on probabilistic explanations of some puzzling experimental findings, as given by Oaksford, Chater, Stevenson, and Over. A presupposition of probability theory is the Boolean event structure over which probability is defined, and van Lambalgen argued that this sits badly with the intensional aspects of human reasoning. Further, certain non-monotonic aspects of reasoning cannot be modelled in a Bayesian way because they necessitate making changes to the probability assignment instead of simply updating on them.

Max Albert (Economics, Giessen) and Jon Williamson (Philosophy, Kent) spoke about probabilistic and statistical methods in general. Albert presented a formal model for scientific progress in which different statistical methodologies, both classical and Bayesian, could be represented as different strategies, and then argued that the classical methodology provides a better foundation for progress. Williamson looked at more recent developments in statistics, discussing the use of logical constraints, causal knowledge, network representations and objective Bayesianism in improving and speeding up statistical procedures.

The general focus of the workshop was on the interplay between, on the one hand, the use of probability theory in modeling human agents and, on the other, the use of statistical models in dealing with data about human agents. But the workshop itself gave an even wider range of topics to think about. It underlined the importance of seeking connections between philosophy and science, but also the complexity of that task and the fun that both philosophers and scientists may derive from it.

Jan-Willem Romeijn
Philosophy, University of Groningen

Call for papers

Spontaneous Generations: A Journal for the History and Philosophy of Science, deadline 15 June 2007.

Combining Probability and Logic, special issue of the Journal of Applied Logic, deadline 1 July 2007.

Special Issue of Foundations of Science: Mathematics and Argumentation, deadline 1 November 2007.

LU: Logica Universalis, Publisher: Birkhäuser Basel

Prize

THE WOLFRAM 2,3 TURING MACHINE RESEARCH PRIZE, in celebration of the 5th anniversary of Stephen Wolfram's A New Kind of Science.

§4
EVENTS

THE SQUARE OF OPPOSITION: Montreux, Switzerland, 1-3 June 2007.

SEMINAR: Robert Tragesser's book Phenomenology and Logic (Cornell University Press, 1977), IHPST (Paris 1/CNRS/ENS), Paris, France, 2 June 2007.

LFCS'07: Symposium on Logical Foundations of Computer Science, CUNY Graduate Center, New York City, June 4 - 7, 2007.

TALK: Royal Institute of Philosophy seminars, University of Sussex, 8th June, Dr Ron Chrisley 'A realist, bivalent semantics for logics that reject the law of excluded middle'.

ARCHE VAGUENESS CONFERENCE: 8-9 June 2007, St. Andrews.

OBAYES6: The 6th International Workshop on Objective Bayesian Methodology, Rome, Italy 8-12 June, 2007.

BAYESIANITY: Foundations and Applications, to be held in Sao Paulo June14 and 15, 2007.

LOGICA 2007: Hejnice Monastery, Czech Republic, 18-22 June 2007.

CIE 2007: Computability in Europe 2007: Computation and Logic in the Real World University of Siena, Siena, 18-23 June 2007.

DYNAMIC LOGIC MONTRÉAL: Université du Québec à Montréal, June 19th to 22nd, 2007.

WORKSHOP: AHRC Project: Transcendental Philosophy and Naturalism, Transcendental Arguments Workshop, Mordan Hall, St. Hugh's College, Oxford, OX2 6LE, 21-22 June 2007.

TARK XI: Eleventh Conference on Theoretical Aspects of Rationality and Knowledge, Brussels, 25-27 June 2007.

1ST GPMR: 1st GPMR Workshop on Logic & Semantics "Medieval Logic and Modern Applied Logic", Rheinische Friedrich-Wilhelms-Universität Bonn, Germany, June 28-30, 2007.

WOLLIC'2007: 14th Workshop on Logic, Language, Information and Computation, Rio de Janeiro, Brazil, July 2-5, 2007.

BSPS: British Society for the Philosophy of Science, Bristol, 5-6 July 2007.

MAXENT 2007: The 27th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, The Saratoga Hotel Saratoga Springs, New York, USA, July 8-13, 2007.

SCIENTIFIC MODELS: SEMANTICS AND ONTOLOGY: Barcelona, July 9-10 2007

LCC'07: Workshop on Logic and Computational Complexity, (affiliated with LICS 2007), Wroclaw, Poland, 15th July 2007.

CADE-21: Workshop on Empirically Successful Automated Reasoning in Large Theories (ESARLT), 15th July 2007.

LFMTP'07: International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice, Affiliated with CADE-21, Bremen, Germany, July 16 2007.

LORI: Logic, Rationality and Interaction, Beijing, 5-9 August 2007.

TANCL'07: Algebraic and topological methods in non-classical logics III, 5-9 August 2007, Oxford.

WORKSHOP: Construction and properties of Bayesian nonparametric regression models, Isaac Newton Institute for Mathematical Sciences, Cambridge, UK, Aug 6-10 2007.

LMPS: 13th International Congress of Logic, Methodology and Philosophy of Science, Beijing, 9-15 August 2007.

UNI-LOG: 2nd World Congress and School on Universal Logic, Xi'an, 16-19 August 2007.

C&O:RR-2007: The Third International Workshop on Contexts and Ontologies: Representation and Reasoning, August 21, 2007, CONTEXT Workshop Program, Roskilde University, Denmark.

LSFA'07: Second Workshop on Logical and Semantic Frameworks, with Applications, August 28th, 2007, Ouro Preto, Minas Gerais, Brazil

ASAI 2007: IX Argentine Symposium on Artificial Intelligence Mar del Plata, Argentina, August 27-28, 2007.

PROGIC 2007

The Third Workshop on Combining Probability and Logic, University of Kent, 5-7 September 2007.

IDA 2007: The 7th International Symposium on Intelligent Data Analysis, Ljubljana, Slovenia, September 6-8, 2007.

DYNAMICS OF KNOWLEDGE AND BELIEF: Workshop at KI-2007, 30th Annual German Conference on Artificial Intelligence, Osnabrück, 10 September 2007.

CSL 2007: Computer Science Logic 11-15 September, 2007, Lausanne (CH).

AIPL-07: Workshop on Artificial Intelligence Planning and Learning, Providence, Rhode Island, September 22, 2007, organized in conjunction with the International Conference on Automated Planning and Scheduling (ICAPS-07).

SPRING BAYES 2007: The 4th annual meeting of Australasian Society for Bayesian Analysis (ASBA) will take place in Coolangatta, 26-28 September, 2007.

TBILISI: The Seventh International TBILISI Symposium on Language, Logic and Computation, 1-5 October 2007.

Reason, Intuition, Objects: The Epistemology and Ontology of Logic, Buffalo, 13 October 2007.

LPAR 2007: Yerevan, Armenia, 15th-19th October 2007.

ECSQARU'07: Ninth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, October 31, November 1-2 2007, Hammamet, Tunisia.

International Conference on Infinity in Logic and Computation: 3-5 November 2007, University of Cape Town, South Africa.

Workshop: 3rd Workshop on Uncertainty Reasoning for the Semantic Web, Busan, Korea November 12, 2007.

EPSA07: 1st Conference of the European Philosophy of Science Association, Madrid, 15-17 November 2007.

Graduate Conference: 1st Cambridge Graduate Conference on the Philosophy of Logic and Mathematics, 19th-20th January 2008 St. John's College, Cambridge.

ISBA08: 9th World Meeting of the International Society for Bayesian Analysis (ISBA), Hamilton Island, Australia, 21st-25th July 2008.

§5
Jobs

Visiting Assistant Professor: Visiting Assistant Professor of Philosophy full-time, University of Hawaii at Hilo, one year only with possibility of renewal, ronald@hawaii.edu, Deadline: First review of applications will begin May 24 and continue until the position is filled.

Lectureship in Philosophy: College Stipendiary Lectureship in Philosophy, St John's College, Oxford, Closing date 6 June 2007.

Lectureship in Logic: Department of Philosophy, University of Auckland, closing date 8 June 2007.

Lectureship in Philosophy: Hull, Closing date: 8 June 2007.

Teaching Fellowship: Philosophy, UCL, from September 17th 2007, closing date for applications Friday, 22nd June 2007.

Postdoc: Natural Language Processing and Machine Learning Post-Docs at The Cognitive Computation Group at the University of Illinois at Urbana/Champaign, rbking@uiuc.edu.

Postdoc: Rensselaer Polytechnic Institute, Human-Level Intelligence, Laboratory at the Rensselaer Department of Cognitive Science, cassin@rpi.edu.

§6
Courses and Studentships

Courses

Master Program in Intelligent Systems: University of Lugano in collaboration with IDSIA, Switzerland, Enrolment deadline: 1 July 2007.

Research Master in Language, Cognition, Action, and Mind Studies: The Institute for Logic, Cognition, Language, and Information of the University of the Basque Country (Donostia-San Sebastian).

Logic Summer School: Italian Association of Logic and its Applications (AILA), Italian Society for Logic and Philosophy of Science (SILFS), Palazzo Feltrinelli, Gargnano, Italy, 26 August - 1 September 2007.

SECEVitA 2007: Summer School in Artificial Life and Evolutionary Computing, 31 August – 4 September 2007, Baia Samuele, Ragusa, Italy

Second Indian Winter School on Logic:, January 14-26, 2008, IIT Kanpur.

Studentships

BSPS Doctoral Scholarship in Philosophy of Science, closing date 1st August 2007.

Acknowledgements

The Reasoner is supported by a Strategic Research Development Fund grant from the University of Kent. It is a development of the progicnet academic network supported by the Leverhulme Trust.