
THE REASONER

VOLUME 2, NUMBER 10
OCTOBER 2008

www.thereasoner.org
ISSN 1757-0522

CONTENTS


- §1 Editorial
- §2 Features
- §3 News
- §4 Introducing ...
- §5 Events
- §6 Jobs
- §7 Courses and Studentships

§1 EDITORIAL

Interdisciplinarity is something that many pay lip service to but few are prepared to act on. Funding bodies claim to be eager to fund interdisciplinary projects but when it comes to making decisions, the peer review system tends to ensure that tradition takes over. And where successful, interdisciplinarity is often a victim of its own success: valuable work across disciplinary boundaries often results in ... a new discipline. According to the guidelines of the European Research Council, for instance, bioinformatics no longer counts as interdisciplinary because it has established itself as an au-

tonomous discipline. Is interdisciplinarity always self-defeating?

Perhaps 'Causistics' will one day emerge as a new discipline studying causal inference in the sciences, just as Statistics emerged to study probabilistic inference in the sciences, but that day is a long way off. In the mean time the conference series on [Causality in the Sciences](#) offers a forum for a variety of disciplines and interests to meet and exchange ideas. The latest incarnation of this series was the [Causality Study Fortnight](#) held recently in Canterbury. It took two days of tutorials, three days of conference and five days of seminars to start setting a common language, one we hope to build on. Our two interviewees this month spoke at the Fortnight and they talk here about causal reasoning.



What's the role of *The Reasoner* in this interdisciplinary enterprise? Let us remind you that the aim of *The Reasoner* is to present a variety of perspectives on topics related to reasoning, inference and method.

We urge those across the spectrum of disciplines to submit features and items of news—it's a good way to promote your research, initiate discussions and generate interest in your area, and it's the fastest way to get to the heart of a scientific and philosophical community with diverse interests and skills.

Please check out <http://www.thereasoner.org> for submission guidelines. We will print articles, items of news, letters, conference announcements, calls for



papers and job announcements.

[Federica Russo](#)
Philosophy, Louvain & Kent

[Jon Williamson](#)
Philosophy, Kent

§2 FEATURES


Interview with Nancy Cartwright

[Nancy Cartwright](#) is Chair of the [Centre for Philosophy of Natural and Social Science](#) and Professor of Philosophy in the [Department of Philosophy, Logic and Scientific Method](#) at the London School of Economics. She is also Professor of Philosophy at the University of California, San Diego. She is known for her work on the philosophy of science, causal inference and evidence-based policy.

Jon Williamson. Could you tell the readers something about your intellectual history?

Nancy Cartwright. At the beginning of my career, when causality was almost a forbidden topic, I wrote ‘Causal Laws and Effective Strategies’. This paper argued that we cannot answer questions of what we should do to bring about the outcomes we want if we do not distinguish between causal laws, which may be backed by probabilistic relations, and the probabilistic relations themselves. About the same time others were beginning to develop causal decision theory. Our views dovetailed nicely. I offered a formula to use for the probability of a causal counterfactual (if C were to occur E would occur as a consequence) to replace the more usual conditional probability in decision theories. My proposal was the now familiar partial conditional probability, $P(E|C \& K_i)$, where K_i is a state description over all the ‘other’ causes at work in the situation in which the counterfactual is evaluated.

Although this formula has long been used in the social sciences and is part of Suppes’s probabilistic theory of causality, I arrived at it by, on the one hand, a complicated analysis of what would happen in individual cases in a population of mixed K_i ’s and, on the other, by my personal discovery of Simpson’s paradox, which was not known in the philosophical literature. (I noticed Simpson’s paradox by analyzing and generalizing Wesley Salmon’s examples where causes can de-



crease the probability of their effects. It was so counterintuitive that a correlation that holds in a population could be reversed in each partition, and vice versa, that I couldn’t believe my calculations. I finally broke down and asked one of Stanford’s great statisticians, Persi Diaconis, where my calculations were going wrong. Persi laughed a bit at me and taught me the name ‘Simpson’s paradox’.)

I have worked on causality on and off ever since, first at Stanford where the bulk of my other work was in philosophy of physics and since coming to LSE as part of my endeavour to learn philosophy of the economic and social sciences. I tell this story about the start of my work on causality because of my answers to some of your other questions.

JW. Why do you think it’s so important to understand causal reasoning?

NC. Ever since ‘Causal Laws and Effective Strategies’ I have thought it essential to understand causality in order to be in a better position ‘to change the world’. But I am now deeply concerned that that is mistake.

JW. In a nutshell, what are your views about causal reasoning?

NC. As with any concept, metaphysics, method and use must be mutually reinforcing. Our account of what causality is and how it behaves should justify both our methods for ascribing the concept and the inferences we wish to make from the concept. I now see that there is no existing concept or theory of causality that does both without merely tacking them together. (Consider $P(E|C \& K_i)$: this is not really the probability that E would result were I to do C given K_i , since it does not take into account the effects of other factors nor new factors I will introduce when I do C or changes I might make to the underlying nomological structure. $P(E|C \& K_i) > P(E|C \& K_i)$ is instead a test for a causal law: C causes/can cause E ; or a measure of the strength of C ’s capacity/power to produce E . So the probabilistic theory of causality directly justifies an important method for finding causes but does not play the role in real-life counterfactual causal prediction that I, along with many causal decision theorists, assigned it.)

Tacking is common now in many accounts of causality, notably some manipulation theorems for Bayes nets, Woodward’s 2-clause characterization of causality and certain kinds of empirical ‘robustness’ tests for causal conclusions. But tacking is not enough. It cannot justify why satisfying the clauses in the characterization of causality that underwrite method is enough to suppose the clauses that justify inferences will be satisfied. So perhaps causality does not underwrite strategy after all, or only causal laws that have a special source. (This is my current best hypothesis, with the favoured sources being causal powers or what I call ‘nomological machines’.) But then, can the same source not also generate non-causal associations that underwrite strat-

egy predictions equally well? If so, causal knowledge is not even special for effective strategies when it has a special source.

These concerns about the usefulness of causal knowledge fit with the severe pluralism I hold. There is no such thing as the causal relation. There are a huge, perhaps, indefinite number of empirical relations: x is to the right of y , x is mother of y , x pushes y , x is bigger than y , x attracts y ,... For various different ordinary purposes, generally having to do with responsibility and management, some of these get labelled 'causal'. In science we study a large number of different kinds of systems of lawlike relations. Many of these lawlike relations get labelled 'causal' as well, again for different vaguely related reasons, in science having to do with prediction, control and explanation.

JW. What are you working on now?

NC. Just this problem about what use knowledge of causal laws can be.

JW. Do you think it's helpful to pay attention to what other disciplines say about causality?

NC. Yes.

JW. Can you recommend a topic or two for the graduate student interested in causal inference and scientific method?

NC. I am very keen that philosophy of science should be of genuine use in the natural, social and policy sciences (and am hence a strong supporter of the Society for the Philosophy of Science in Practice, which I urge you all to join). So I would like to direct students of causality to work seriously on this problem of how to put causal knowledge to use, both in very controlled environments and in much more rough and tumble life where we also seem to be able to get around better with the aid of our causal knowledge? How is this possible and how can we do it better?

Interview with Jim Joyce

[Jim Joyce](#) is Chair and Professor of Philosophy and Statistics at the University of Michigan. He is best known for his work on causal decision theory.

Federica Russo. Thank you so much, Jim, for giving us an interview this month. Perhaps you can tell our Reasoners about your background and what brought to you decision theory.

Jim Joyce. It was an accident. In the late 1980s I was working on a Ph.D. dissertation on the role of randomness assumptions in statistical physics. In the course of this I became interested confirmation theory and the question of how one might go about testing these assumptions. The first philosophical litera-



ture I read on the topic, which involved broadly syntactical approaches to confirmation, left me frustrated. It seemed clear that the 19th-century picture of scientific testing as essentially deductive had long been superseded by a picture in which probabilistic approaches should play the leading role. Now, at the time I knew nothing of Bayesianism (I may have heard the word). As an undergraduate I pursued math and philosophy and had taken three courses in statistics, but all taught in the frequentist style. So I knew nothing of "priors," "posteriors," "Bayes factors," and the like. My first exposure to Bayesian ideas came from looking at some work of E. T. Jaynes, an "objective Bayesian," and from reading Clark Glymour's *Theory and Evidence*. I found Jaynes's interesting, but was put off by the a priori nature of his approach. Glymour, on the other hand, made a big impression, especially his attempt to refute Bayesian confirmation theory via the "problem of old evidence," which initially seemed convincing to me (though it no longer does). Sometime after Glymour, I got hold of a copy of Leonard Savage's magnificent *Foundations of Statistics*, not really knowing what it was but hoping to get a textbook on foundational aspects of frequentist estimation and hypothesis testing. Savage's book is, of course, the masterwork of Bayesian statistics and of modern expected utility theory. Reading it left me a committed Bayesian and it sparked in me a lasting interest in decision theory.

A number of aspects of Savage's framework interested me: his behaviorist interpretation of beliefs and desires (which I did not like); his justification of probabilistic coherence for degrees of belief using representation theorems (which I did like); the structural assumptions he used to secure the uniqueness of probability representations (which I did not like); and the problem of "small works" (which I found fascinating). Deciding to think about these issues, I changed my thesis topic and began to think about decision making. I subsequently read works by F. P. Ramsey ("Truth and Probability"), Bruno de Finetti (*Theory of Probability*), Skyrms (*Causal Necessity*), Jeffrey (*The Logic of Decision*) and Gibbard and Harper ("Counterfactuals and two kinds of Expected Utility"), all things I recommend unreservedly to others. My dissertation ended up being a critique of what I saw as shortcomings in Savage's approach, and the development of a version of "causal" decision theory had some of the advantages of Jeffrey's "evidential" approach without its limitations. This work ultimately led to *The Foundations of Causal Decision Theory* (1999).

FR. As you explained in your tutorials at the [Causality Study Fortnight](#), decision theory (DT) is a formal account of practical rationality and tells us under what conditions an agent should act. I understand you take decision theory to have a crucial normative component. But how big is the gap between decision theory and ex-

perimental studies in psychology on how people choose and act, then?

JJ. That's a big question! The gap is huge—real people are nothing like expected utility maximizers—but the real question is what we should do with this information. My opinion is that there is no uniformly correct way to interpret the empirical data. Some of patterns of decisions that people make, and the justifications they offer for them, are just irrational, while others require us to rethink foundational aspects of decision theory. For a (controversial) example of the first consider the famous Allais paradox. My view is that the rationale for the standard observed choices in Allais reveals that people have beliefs about which gambles are riskier than others that are simply incoherent. There is, in my opinion, no consistent notion of “risk” that can play the role that Allais choosers want it to play. I'd say the same sorts of things about “faming effect” cases in which a person's preferences depend on the ways in which alternatives happen to be described. On the other hand, I think that the usual pattern of preferences in the Ellsberg paradox requires us to recognize that (a) decision theory cannot require agents to have sharp probabilities and utilities, (b) that a realistic decision theory will, instead, represent agents using sets of probability/utility pairs, and (c) in decisions like Ellsberg, we cannot interpret an agent's forced choices as revealing much of anything about her preferences since an agent with indeterminate preferences can still choose things. In still other cases, the right thing to do might be to see how much of the empirical data we can accommodate within standard expected utility theory by thinking a bit more deeply about how people's desires work. An excellent example of this sort of work is “[A Model of Reference-Dependent Preferences](#)” by the Berkeley economists Botond Koszegi and Matthew Rabin. Koszegi and Rabin develop a model in which the well-known phenomenon of loss aversion is portrayed as resulting from a standard utility and an externally determined “reference point.” I do not agree with Koszegi and Rabin on some important points, but I do think more decision theorists should be pursuing work of this kind.

FR. You are among the staunchest proponents of *causal* decision theory. Perhaps some of our Reasoners are not entirely familiar with this variant of DT. Would you be able to give them a nutshell-explanation of the distinctive feature of causal DT? Why, in your view, can't we get by without causation?

JJ. Early formulations of decision theory did not incorporate causal information, but all of them endorsed some version of the dominance principle. Dominance reasoning says, roughly, that if one act produces a better outcome in every state of the world, then that act is to be preferred. This sort of reasoning is only valid, however, when acts cannot alter the probabilities of states. Consider a highly allergic man who carries an

epinephrine kit because he is likely to die of anaphylactic shock if stung by a bee. Once stung, the man must decide whether to have a mildly painful injection that is certain to save him or to forgo the shot and hope anaphylaxis does not occur. Such a person would be crazy to reason this way: either I go into shock or I don't, either way I would rather avoid the pain of the shot, so I won't take the shot. This inference ignores the fact that the decision to take the shot is highly causally relevant to the occurrence of anaphylactic shock. Moral: Dominance does not apply when acts can causally influence states of the world on which their consequences depend. The question is how to incorporate this insight into decision theory. One option is to insist that dominance should only be invoked when states are *independent* of acts. Alternately, one can replace the usual unconditional form of the expected utility law $EU(act) = \sum Probability(state) \cdot utility(act \& state)$ with a requirement to maximize conditional utility $EU(act) = \sum Prob(state \text{ given } act) \cdot util(act \& state)$, with the understanding that the $Prob(state \text{ given } act)$ captures the degree to which performing the act will alter the state's probability. Here the value of an act is obtained by weighting the utilities of each outcome by the probability of the state that brings that outcome about where this probability is conditioned on the act in question being performed. Dominance can then be replaced by a principle that requires one to prefer act *A* to act *B* when $Prob(s \text{ given } A) \cdot util(A \& s) > Prob(s \text{ given } B) \cdot util(A \& s)$ for all states *s*.

On either approach rational decision making requires having views about the dependence relationships that obtain between actions and states. A central question is how to represent these views. Here theorists divide into two camps. Evidentialists, like Richard Jeffrey (at one time), advised agents to choose acts that provide evidence for thinking that desirable results will obtain (even when these acts do not causally promote those results). This amounts to reading the “given” as ordinary conditional probabilities, $Prob(s \text{ given } A) = Prob(s \text{ given } A) / Prob(s \text{ given } A)$. Causalists advise agents to choose actions that causally promote desirable outcomes (even if these acts provide evidence for undesirable outcomes they do not promote). This amounts to reading the “given” as a special kind of causal probability, which I write as $Prob_A(s)$, that measure the degree to which *A* is a promoting cause of *s*. (Note: these are not ordinary conditional probabilities.)

Now, much of the discussion about the evidential and causal approaches has, unfortunately, focused on Newcomb problems, but I want to emphasize that the fundamental difference between evidential and causal decision theories concerns, not what agents should do, but whether we can adequately characterize their rationale for doing what they should do without making explicit reference to their beliefs about what their choices are

likely to cause. Causal decision theorists maintain that there is no avoiding causality. To know what an agent should do in a given decision situation we must know what she believes about the *effects* of her actions, and this is only possible if we make reference either to her beliefs about propositions with explicitly causal content or to forms of belief revision that are subject to explicitly causal constraints. Evidential theorists dispute this. They claim that, for the purposes of doing decision theory, we can capture the relevant beliefs about causes and effects by appealing to nothing more than the agent's ordinary subjective probabilities for non-causal propositions.

I think that the causal decision theorists are right. While there are lots of hard questions to be answered about how best to represent causal reasoning (see the [slides](#) from my presentation for some relevant information), I do not see how CDT could be wrong. When one acts deliberately, as an *agent*, one is aiming to change the world for the better. Actions are causes, pure and simple, and what we need to know when making a decision is what act is likely to cause the best results. The evidence that the act provides about good or bad things is a side effect, not the main show.

FR. What do you think has been the most relevant improvement in decision theory after Savage's formulation of DT? Is it something related to the *foundations* or to the technicalities of DT?

JJ. That's a hard one. Without committing myself to saying what improvement is "most" relevant, I would say that one major, major advance is the idea (I am not sure who had it first) that decision theorists need to think about the way that agents deliberate themselves into states of a decision-theoretic equilibrium. Ken Binmore's "Modeling Rational Players I and II" uses this approach to great effect, as does Brian Skyrms in his *Dynamics of Rational Deliberation*. A second major development, due to many people, is the development of a decision theory that allows for agents with imprecise beliefs and desires. P. Walley's *Statistical Reasoning with Imprecise Probabilities* is an excellent reference here.

FR. Decision theory and the debate between causalists and evidentialists is a very fertile area. But which directions is cutting edge research going in at the moment? What do you think are the most exciting problems to be solved in decision theory?

JJ. I guess I do not anticipate major advances at the deep theoretical level, most advances will come in the form of extensions of, or additions to, the decision theories we have. There are two that come most immediately to mind. The introduction of evolutionary methods in game theory has transformed the field, and the idea of thinking about decision making from that perspective has been very fruitful. The works of Brian Skyrms, Ken Binmore, Jason McKenzie, Peter Vanderschraaf

and Patrick Grim are especially worth reading in this area. Second, I think the remarkable recent advances in our understanding of causal reasoning, of which readers of *The Reasoner* are well aware, are beginning to find its way into questions of practical decision-making. There is lots of interesting stuff being done in this area, a lot of it by computer scientists.

FR. Jim, one last question. I am sure you have been asked this a thousand times at least . . . but do you really compute expected utilities when you make decisions? Don't you think sometimes letting fate decide by tossing a coin is the best decision?

JJ. I never compute expected utilities except in the (very rare) situations when I play poker. Then I do try to think in terms of expected values, but I must not be very good at it since I tend to lose (which is why, in an effort to maximize my expected utility, I try to avoid poker whenever possible). For most decisions one does not need fancy decision theory: the answer becomes obvious after a little thinking. Sometimes in major decisions that are especially hard I try to, e.g., see whether the sure-thing principle might apply, but that's about it. In any event, I've never seen expected utility theory as a theory of how people should go about *making* decisions. It is, rather, a tool of assessment that allows us to say how well people succeed in making decisions that are rationalized by their beliefs and values, however they choose to make them.

A Remark About Essential Indexicals

Among many others, Castañeda (1967: Indicators and Quasi-Indicators, *American Philosophical Quarterly* 4, 85–100) and Perry (1979: The Problem of the Essential Indexical, *Noûs* 13, 3–21) have defended the thesis that certain essential indexicals like *I* and *now* are irreducible in thinking: they cannot be substituted by other expressions (in thinking) without losing power in explaining an agent's actions. I will argue that there are at least two distinct ways to understand this thesis and that one of them is too strong. For simplicity, only *I* shall be considered in what follows.

Here is the structure of typical arguments for the irreducibility of *I*:

(IRI) *I* is irreducible in thinking, because for any kind of condition ϕ that is supposed to uniquely identify a person *P* thinking an *I*-thought, *P* might not realize that she herself is the ϕ -er.

In this scheme, *to realize* is linked to belief. From the fact that *S* believes *de re* that $a \neq b$ it follows that *S* doesn't realize that $a = b$, but not the converse. This relationship captures the non-factive aspect of *to realize* that is relevant for the behavior of an agent that doesn't realize that he himself is the ϕ -er, whereas the

factive aspect can be expressed as the condition that he is in fact identical to the ϕ -er. Standard rational belief based on modal logic KD45 needs to be used to express the relationship between not realizing something and belief, because more elaborate notions—such as *de se* belief modelled using structured propositions, property ascription, modes of presentation, or similar technical means—already presume one or another version of IRI.

In a framework similar to the well-known Logic of Demonstratives by Kaplan (1989: Demonstratives, In Almog/Perry/Wettstein: Themes from Kaplan, 481-564), I will now formulate four tentative versions of IRI. A function $\|.\|^g(c)(w)$ evaluates formulas and terms under assignment g with respect to a context c and a world w , where features of contexts (agent viz. speaker, world, time, and addressee) are retrieved by mapping functions with mnemonic names such as $speaker(c)$ for the speaker of a context. Let c^w be the context that is exactly like c except that $world(c) = w$. It is crucial for the argument schemes below that $\|I^*\|^g(c)(w) = speaker(c^w)$, i.e. unlike English I the term I^* is not rigid. (Alternatively, a diagonalization operator may be applied to an ordinary rigid analysis of I .) Corresponding to Kaplan's *Dthat* operator on terms, a sentential operator $@$ may be defined as $\|@A\|^g(c)(w) := \|A\|^g(c)(world(c))$, and the following quantifier will be used instead of a iota operator: $\| \iota x(A)B \|^g(c)(w)$ is true if there is exactly one x -variant h of g such that $\|A\|^h(c)(w)$ is true, and for this assignment $\|B\|^h(c)(w)$ is true; false otherwise. For finitely many speakers S , let $\|\mathcal{B}_x A\|^g(c)(w)$ be true if $g(x) \in S$ and in all worlds w' compatible with what x believes, i.e. all w' such that $R_{\mathcal{B}}(w, w', g(x))$, $\|A\|^g(c)(w')$ is true; false otherwise. $R_{\mathcal{B}}$ is serial, transitive and Euclidean with respect to its first two argument places. Finally, E is the existence predicate A a meta-variable for any non-trivial unary predicate.

Given all that, if IRI holds in general, then one or more of the following schemes ought to be valid, i.e. true with respect to any context, world pair in a proper logic of essential indexicals:

- (1) $\forall x[(x=I^*) \supset \diamond \mathcal{B}_{xty}[@(Ay \wedge Ey)] y \neq I^*]$
- (2) $\forall x[(x=I^*) \supset \diamond \mathcal{B}_{xty}[Ay \wedge Ey] y \neq I^*]$
- (3) $\forall x[(x=I^*) \supset \diamond \mathcal{B}_{xty}[@(Ay \wedge Ey)] y \neq x]$
- (4) $\forall x[(x=I^*) \supset \diamond \mathcal{B}_{xty}[Ay \wedge Ey] y \neq x]$

Which scheme is the right one? Consider (3) and (4) first. In both of them, the referent of I^* is determined externally to the agent's own beliefs, i.e. from a 3rd-person perspective. If an occurrence of I^* outside the scope of a belief operator is understood as the formal analogue to an ordinary use of English I , the bound variable x in these schemes represents ordinary, non-essential readings of I from a 3rd-person perspective,

though formalized in an unusual way using a non-rigid, first-person indexical and applying the wide scope theory to it. Hence, the schemes are unsuitable for expressing IRI, since IRI relates beliefs about oneself formed on the basis of *first-person* thoughts with those formed on the basis of 3rd-person means of identification.

What about the other schemes then? Scheme (1) expresses an implausibly strong condition. It roughly says that for any property A in any context the agent (viz. the speaker or thinker) of that context might believe that the individual object that actually exists and actually satisfies A is not identical to himself. In other words, no property might actually uniquely identify something which actually exists and of which the speaker thinks as himself. This would make I in thinking epistemically irreducible and a close relative of qualia. Yet it seems that the typical examples by Perry (1979) and others don't support this view. After all, if John Perry doesn't recognize himself as the only person with such-and-such properties (say, in the supermarket), only whatever he believes to have these properties can be relevant for his behavior and not whatever is *actually* uniquely determined by them or not.

This leaves us with the much weaker scheme (2). It roughly says that for any property A in any context it is possible that the agent of the context believes that he himself is not identical with the object that exists and uniquely satisfies A according to his beliefs. Still, there might actually be a (presumably complex) property that uniquely determines what the agent believes to be himself. This scheme merely contrasts the possibility of error of identifying something by means of certain properties with the infallibility of 'indexical reference' in thinking. But to this it can be replied that whenever somebody has an I -thought, i.e. he thinks about himself, he trivially has an I -thought and not an A -thought, since the I -thought is defined as the one linked to the respective I -behavior (Ah! It is myself who pours sugar on the floor! So I'll clean it up.) Any A -thought linked to the same behavior would invariably be an I -thought. In this sense, (2) is harmless and essential indexicals might be less problematic to a reductionist view than previously thought.

Erich Rast
Philosophy, Lisbon

This is Nonsense

In his *Paradoxes* (1995: CUP, 149) Mark Sainsbury presents the following pair of sentences:

- Line 1: The sentence written on Line 1 is nonsense.
Line 2: The sentence written on Line 1 is nonsense.

Sainsbury (1995: 149, 154) here makes three assertions: (1) The sentence in Line 1 is so viciously self-referential that it falls into the truth-value gap. The sentence is really nonsense. (2) The sentence in Line 2 is by contrast true. For it states precisely that the sentence in Line 1 is nonsense. (3) The two sentences in Lines 1 and 2 are an example of the principle that two sentence tokens of the same sentence-type can have different truth-values, although they have the same reference and state the same property of the object of reference.

Sainsbury's assumptions are false in all three cases. In order to demonstrate this, let us ask what would be a suitable definition of 'nonsense'. A minimal condition for sentences would be the following: nonsensical sentences express no proposition. Since a sentence which expresses no proposition cannot have a truth-value, it follows that a nonsensical sentence does not have a truth-value. Nonsensical sentences are thus neither true nor false.

Let us suppose that the sentence in Line 1 is true. In this case what the sentence in Line 1 states is true and thus the sentence in Line 1 is nonsense. If the sentence is nonsense, then it expresses no proposition and so has no truth-value. If it has no truth-value, then it is at least also not true. The assumption that the sentence in Line 1 is true thus leads to a formal contradiction: the sentence in Line 1 is both true and untrue. Thus with the help of the *reductio ad absurdum* it follows that the negation of the original assumption is true: the sentence in Line 1 is not true. If it is not true, then what it states is not true. Since it states that it is nonsense, it is therefore in fact not nonsense. The falseness of the sentence in Line 1 does not lead to a contradiction. A sentence may be false and at the same time sensible. The sentence in Line 1 is therefore a sensible, though false, sentence. It is, however, not a contingently false sentence. For it is impossible for it to state truly about itself that it is nonsensical. If, however, the sentence in Line 1 is false, the sentence in Line 2 is not true. For, as has been demonstrated, the sentence in Line 1 is not nonsense, whereas the sentence in Line 2 states the opposite. The sentence in Line 2 is therefore also false. This means that both sentences are both sensible and false.

Thus, Sainsbury's first two assumptions are false (the sentence in Line 1 has a truth-value and the sentence in line 2 is not true). Their falsity also makes his third assumption false. The self-referential nonsense sentence in Line 1 and its related twin in Line 2 are not the only examples of their types. In fact, the following sentence pairs belong to the same type:

No truth conditions

Line 3: The sentence in Line 3 has no truth conditions.

Line 4: The sentence in Line 3 has no truth conditions.

On condition that a sentence is true if and only if

its truth conditions are fulfilled and false if and only if its truth conditions are not fulfilled, a sentence which does not have any truth conditions at all is neither true nor false. Let us suppose the sentence in Line 3 is true. Under the given conditions this sentence would then be neither true nor false and the result is once more a *reductio ad absurdum* as in the nonsense sentence in Line 1. Thus the sentence in Line 3 is false and has truth conditions, and the identically formulated sentence in Line 4 is also false.

Truth-value gap

Line 5: The sentence in Line 5 falls into the truth-value gap.

Line 6: The sentence in Line 5 falls into the truth-value gap.

Sentences which fall into the truth-value gap have *per definitionem* no truth value and are accordingly neither true nor false. If sentence 5 wished in truth to state about itself that it falls into the truth-value gap, it would be neither true nor false. The result would again be a *reductio ad absurdum*. Thus the sentence in Line 5 is false and does not fall into the truth-value gap. The identically formulated sentence in Line 6 is then also false.

The self-referential nonsense sentences and their related twins have in common the fact that they can all be derived from the following sentence pair, which in turn defines their common nature:

No truth-values

Line n: The sentence in Line n is not true and not false.

Line n+1: The sentence in Line n is not true and not false.

If the sentence in Line n were true, it would be not true and not false and hence at least not true. The *reductio ad absurdum* leads to the fact that the sentence in Line n is not true. If it is not true, then it is not the case that it is not true and not false. Therefore it is either true or false. Since, as we have already seen, it is not true, it can—in accordance with disjunctive syllogism—only be false. It being false is consistent with being either true or false. If, however, the sentence in Line n is false, then what the sentence in Line n+1 states is not the case, so that the sentence in Line n+1 is then also false.

It is thus completely impossible for a sentence *A* to state of itself in truth that it is nonsense, that it has no truth conditions, that it falls into a truth-value gap or, putting it briefly, that it has no truth-values. It is also impossible for a sentence token *B* of this sentence-type to be true as long as it refers to the corresponding self-referential sentence *A* and makes the same statement about *A* as *A* does about itself. Nonsense sentences *A* and *B* of this kind cannot be examples of the principle that two sentence tokens of the same sentence-type can

have different truth-values, although they have the same reference and state the same property of the object of reference.

(My thanks to Mark Sainsbury, Francesco Berto, and Dominic Kaegi for helpful comments.)

Gregor Damschen
Philosophy, University of Halle & University of
Lucerne

§3 NEWS

Foundations of Logical Consequence

A new project researching into the Foundations of Logical Consequence will start at the Arché Centre for Logic, Language, Metaphysics and Epistemology at the University of St. Andrews in January 2009. It is funded by the Arts and Humanities Research Council.

Two approaches have dominated discussion of logical consequence in recent years, the model-theoretic and the inferentialist. The model-theoretic analysis identifies logical consequence with truth-preservation in models: every model of the premises must also be a model of the conclusion. Such models can, in Etchemendy's terminology, be either interpretational (varying the interpretation of the vocabulary) or representational (varying the "facts"). In contrast, the inferentialist analysis of consequence concentrates on the notion of proof or derivation, consisting in the application of a set of rules of inference. Rather than judge the rules as correct if they are truth-preserving over models, the inferentialist approach takes the rules as autonomous, constitutive of the meaning of at least the logical terms they contain. For example, the reason Modus Ponens (to infer B from A and 'if A then B') is a correct form of inference is not because it preserves truth; on the contrary, 'if' gains its meaning from being that expression which permits inferences of this form. The order of explanation is reversed.

The early stages of the project will consider and contrast these two approaches and seek clarity over their statement. But what is most important about the debate over the correctness of these two analyses of the concept of logical consequence is how successful each of them is at explicating and critically assessing our inferential practice. The concept may be a theoretical one, but it is one with important practical implications, for inferential behaviour is universal and an essential component of any rational practice. Those practices throw up theoretical questions. For example, can logical theory be revisionary of our actual practice? Should logic be topic-neutral, or can there be a variety of logics, each

equally legitimate but suitable for different purposes or different areas? Is all valid inference formal, preserved through uniform substitution for non-logical terms or concepts, or are there material inferences, valid partly in virtue of the non-logical words used? Is there a clear criterion by which to distinguish logical terms and concepts from non-logical and descriptive terms? Can a systematic logic be developed for semantically context-sensitive expressions, like demonstratives and tenses? What is the connection between logical consequence and correct inference? How is tacit knowledge of logical principles possible, given that most actual inferential practice is carried on without reference to its theoretical underpinnings?

The driving concern of the project on the foundations of logical consequence will be to evaluate how successful each of the two approaches, model-theoretic and inferentialist, are in response to these and other questions. Further details of the project can be found on the [AHRC website](#) and on the [Arché website](#). Funding for the project includes two Post-Doctorial Fellowships and two Postgraduate Studentships.

Stephen Read
Philosophy, St. Andrews

Induction: Historical and Contemporary Perspectives, 8–10 July

The aim of the international conference "Induction: Historical and Contemporary Perspectives", which took place on 8-10 July 2008 at Ghent University, was to address the notion of induction and inductive reasoning from historical, philosophical and formal perspectives. In the first plenary lecture Laura J. Snyder (St. John's University) offered a new—anti-hypothetico-deductive—reading of John Herschel's "Preliminary Discourse on the Study of Natural Philosophy" (1830). She argued that Herschel believed that analogical reasoning was a key part of scientific discovery and that the scattered comments Herschel made about "bold leaps" are meant to refer to instances of analogical inference, not conjectures or guesswork. Two plenary lectures either directly or indirectly addressed Goodman's "grue paradox". While Samir Okasha (Bristol University) defended F. Jackson's solution to Goodman's supposed paradox, according to which there is no "grue paradox" and that there is no such thing as "projectibility", James W. McAllister (University of Leiden) drew philosophical lessons from a quantitative variant of Goodman's riddle of induction: as data sets can in principle be decomposed into a pattern and a remainder or noise term in any one of infinitely many different ways, different data patterns, which provide evidence that different structures in the world are available, will be brought into

focus as our interests change. Steffen Ducheyne (Ghent University) argued in his plenary lecture that William Whewell's tidal research was vital to the development of his philosophy of science. "Tidology" offered him a concrete means to develop and refine his methodological views: the "Special Methods of Induction applicable to Quantity" were the methodological outcome of his research and his defence of equilibrium theory was not a matter of idiosyncrasy but an obvious choice when taking his methodological views with regards to theory-testing into account. John D. Norton (University of Pittsburgh) argued that despite the clear merits of Bayesian Confirmation Theory (BCT), Bayesianism is but one of several useful instruments for assessing inductive relations. Norton offered an overview of the current critiques of Bayesianism and the shortcomings with BCT identified by Bayesians. On the last day a symposium honouring Peter Lipton's work took place. In the first plenary in this symposium, Stathis Psillos (University of Athens) shared his views on the possible harmonisation of Inference to the Best Explanation and Bayesianism. In the second plenary Erik Weber (Ghent University) took stock of Lipton's contributions to scientific explanation and causation. In a second symposium dedicated to Newton's methodology, David M. Miller, Ori Belkind and Eric Schliesser discussed Newton's deductions from phenomena and the role of the rules of philosophizing in his establishment of universal gravitation. The remaining 16 contributed papers discussed a variety of subjects: adaptive logics for inductive reasoning, David Hume, Richard Whately, abduction, Karl Popper, under-determination, simplicity, IBE, Bayesianism, etc. For the full programme see: <http://logica.ugent.be/induction>. During the conference there was excellent interaction between scholars that are working on the problem of induction from different angles—a result which was envisioned by the conference organizers and realized by the participants.

Steffen Ducheyne
Philosophy, Ghent

Computational Social Choice, 3–5 September

On 3–5 September 2008 the 2nd International Workshop on Computational Social Choice (COMSOC-2008) was held at the University of Liverpool. Computational social choice addresses questions of a computational nature in social choice theory, the study of mechanisms for collective decision making, and explores applications of concepts from social choice theory in computer science. It brings together ideas and techniques from a wide range of scientific disciplines, includ-

ing theoretical computer science, artificial intelligence, logic, political science, mathematical economics, and philosophy.

COMSOC-2008 was attended by around 80 participants from over 20 different countries. The programme consisted of five invited talks and the presentation of 36 contributed papers, selected from 55 submissions.

The day immediately preceding the workshop was devoted to introductory tutorials. The day started with a general overview of the various research directions pursued within the COMSOC community, presented by the author of this report. This was followed by in-depth tutorials on two important subfields. Jörg Rothe of the University of Düsseldorf gave an introduction to computational complexity theory, specifically aimed at social choice theorists, illustrating a range of complexity classes with problems naturally arising in social choice, such as the manipulation problem in voting, the computation of power indices, or the solution of fair division problems. Christian List of the London School of Economics gave an introduction to the field of judgement aggregation, which studies the problem of producing a consistent judgement regarding a set of logically inter-related propositions by a group given the individual judgements of the members of that group.

During the workshop itself, invited talks were delivered by Moshe Tennenholtz (Technion), William Thomson (University of Rochester), Tuomas Sandholm (Carnegie Mellon University), Salvador Barberà (Universitat Autònoma de Barcelona), and Rohit Parikh (City University of New York). Moshe Tennenholtz gave an overview of his work on ranking systems, systems where the set of voters and the set of alternatives they vote for coincide. Typical applications are search engines, with webpages being the alternatives and a link from one page to another counting as a vote for that page. William Thomson gave an introduction to the field of fair division, and specifically to the problem of dividing an endowment between a group of claimants when the sum of the claims exceeds the endowment available. Tuomas Sandholm reported on recent work on combinatorial auctions, including the design of highly expressive languages for modelling preferences and requirements of business partners, stressing the relationship between the expressiveness of a mechanism (such as an auction) and its economic efficiency. Salvador Barberà gave an introduction to strategy-proofness of voting rules over restricted domains. While, in general, any voting procedure is manipulable (that is, it will sometimes reward voters who misrepresent their preferences), this problem can be circumvented by restricting attention to voters whose preferences satisfy certain constraints. Finally, Rohit Parikh discussed several applications of epistemic logic to the analysis of social procedures. Other topics covered by contributions to the workshop included belief merging,

the study of tournaments, coalitional voting games, and matching theory.

The proceedings of COMSOC-2008 are available at the [workshop website](#). The next edition of the workshop is planned for autumn 2010 and bids for hosting COMSOC-2010 are currently being solicited.

Readers interested in the field are encouraged to subscribe to the [COMSOC mailing list](#).

Ulle Endriss

ILLC, University of Amsterdam

Phlox, 3–5 September

At the end of last year the DFG-funded research group [Phlox \(Philosophy and Logic of Explanation\)](#) was born at the Humboldt University of Berlin. From the 3rd to the 5th of September this year, a launch workshop on Current Issues in Metaphysics and the Philosophy of Language took place in Berlin to celebrate this event. It brought together young researchers from some of the leading philosophy departments and research groups in Europe. Since only three talks were scheduled per day, there was ample time for in-depth discussion of the various contributions, helped by the fact that all papers were made available in advance.

Katalin Farkas (Budapest) started off the conference by reconsidering the question of whether we can make sense of contingent identity. In opposition to Kripke, she argued for a tentative ‘yes’. Ofra Magidor (Oxford) criticized Williamson’s claim that meta-linguistic safety principles help to explain why we cannot know the alleged sharp cut-off points of vague expressions. She concluded that an epistemicist about vagueness is in need of a better explanation and sketched various options.

Fabrice Correia (Geneva) and Sven Rosenkranz (Berlin) presented an exhaustive classification of mutually exclusive A-theories of time and proposed a non-standard version which combines many of the virtues of B-theories while avoiding many of the vices that afflict more standard A-theories. Elia Zardini (St. Andrews) investigated the semantics and logic of a particular class of modalities, obeying an obliterative principle with respect to a (possibly different) obliterating modality. A formal Kripke-style semantics was developed in detail.

Nick Haverkamp (Berlin) developed a formal framework which brings out the distinction between infinitely improbable and impossible events. He showed that any standard probability function which fudges the distinction can be ‘regularized’ to honour it.

Moritz Schulz (Berlin) proposed a meta-linguistic resolution of a puzzle concerning ‘actually’ sentences and objective chance. He argued that, though it is never a chancy matter whether the propositions expressed by

‘actually’ sentences are true, it is often a matter of chance which propositions are expressed by such sentences.

Stephen Barker (Nottingham) gave advice on how to be a global expressivist. He defended the view that belief has no explanatory role in characterising the nature of assertion. Rather, assertions are expressions of belief only in the sense that they are manifestations of belief, where beliefs are, partly, dispositions to sincerely, and clear-headedly assert.

Dan López De Sa (Barcelona) distinguished two roles of elements in a ‘circumstance of evaluation’: being features of the context shiftable by an operator of the language vs. being features relative to which the objects of attitudes are true. He argued that this distinction allows for a proper taxonomy of positions in the relevant debate.

Benjamin Schnieder (Berlin) took issue with the view that semantic antinomies such as the property-variant of Russell’s paradox prove natural languages to be inconsistent. He showed how we can explain the fact that every speaker runs into the paradox without positing any inconsistency in property-talk.

Miguel Hoeltje

Phlox research group, Humboldt University, Berlin

Causality Study Fortnight, 8–19 September

The [Causality Study Fortnight](#) took place at the Centre for Reasoning (University of Kent) 8–19 September. A major goal of CSF was to bring together philosophers and scientists from a broad range of disciplines to discuss the topic of causality for fifteen long days. The organizers Federica Russo and Jon Williamson managed to host in Kent some forty scholars and to set up the conditions for very fruitful academic exchanges.

Not only did the participants come from very diverse backgrounds, but also from different countries, which made the Fortnight a truly interdisciplinary and international event. Two-day tutorials opened CSF. Causality was broached from different disciplinary perspectives. From Hume to causal pluralism, Julian Reiss gave an efficient overview of philosophy of causality. Kevin Korb provided an introduction to causal modelling and causal discovery that proved very useful for philosophers (and very helpful for understanding some of the talks later during the Fortnight as well). Jim Joyce made it clear that causality matters in decision theory, but he also showed that causal decision theory is insensitive to the choice of one’s particular theory of causality. Finally, David Lagnado presented results and trends in experimental psychology, concerning in particular acquisition of causal knowledge.

The tutorials were followed by a three-day interna-

tional conference: CAPITS 2008. Ten invited talks (Nancy Cartwright, Damien Fennell, Jim Joyce, Kevin Korb, David Lagnado, Michel Mouchart, Stathis Psillos, Miklos Redei, Julian Reiss, Paolo Vineis) and 22 contributed talks were given. Space doesn't allow reporting on all of them, therefore I will just mention the most recurring topics, which works from different perspectives converged on: the role of probabilistic and mechanistic evidence for causal attribution, the causal content of formal probabilistic models, the concept of causality that is at work in various special sciences, the relationship between different levels of causality, the place and role of mechanisms in the metaphysics and in the epistemology of causality, etc.

The third component of the Fortnight consisted of five series of advanced seminars. Levels of causation and mechanisms were discussed in depth by Federica Russo and Phyllis McKay respectively. Alex Freitas and David Corfield covered methodological and philosophical issues related to the [Causality Challenge](#) in the field of automated causal discovery. Philosophy of mind and history of philosophy were not forgotten: Julia Taney discussed mental causation and Ken Westphal dealt with Kant, Newton and Hegel on causation. These advanced seminars triggered fascinating discussions and suggested new paths of research.

Two follow-ups are already scheduled within the [Causality in the Sciences](#) conference series: 'Mechanisms and Causality in the Sciences' to be held in Canterbury in September 2009, and 'Causality and ??? in the Sciences' in Rotterdam in 2010. Abstracts, slides, participants list, and pictures are all available on the [CSF website](#).

Isabelle Drouet
Philosophy, Louvain

Logic of Change, Change of Logic, 10–14 September

Over the past few years, there have been many developments in research on attitude change: from the recent development of a Dynamic Epistemic Logic approach to belief revision to the new interest in the problem of preference change, be it in the modal logic, AGM belief revision or probabilistic frameworks. As the different paradigms, many in different disciplines, begin to camp out their positions, it is perhaps time for us to take a moment to ask ourselves: what do we want from a theory of attitude change? This was the question motivating this year's Prague International Colloquium, Logic of Change, Change of Logic.

Organised by the Institute of Philosophy of the Czech Academy of Sciences, the colloquium assembled philosophers, logicians, computer scientists and

economists, working with a range of frameworks (AGM belief revision, Dynamic Epistemic Logic, Bayesianism, Input-Output Logic, the algebraic approach). The twenty-eight presentations tackled a range of issues: not only belief change and preference change, but also the problem of norm change, the connection between belief revision and conditionals, the relation between desires and beliefs, the role of questions in belief revision and the modelling of intentions. There is no way to do justice to the range of talks. The invited contributions might perhaps give a glimpse of the colloquium.

Two of the invited speakers concerned themselves with general, foundational issues. David Makinson (*Some Design Options for Formal Systems of Cognitive Change*) opened the colloquium with a critical overview of some of the choices which have been made and discussed in the AGM theories of belief revision, and which one might have to come back to when adapting the framework to deal with changes in other sorts of attitudes. In his talk he anticipated some of the issues which were to be taken up again in the conference, such as the correct representation of the state of belief, and the way to model the trigger for change. Sven Ove Hansson (*Can Preference Change Be Modelled after Belief Change?*) opened a day dedicated to preference change with a discussion of the place of preferences with respect to related concepts, such as value, norm and choice, some thoughts about which aspects of belief change carry over to preference change, and finally some proposals for a formal theory of preference change.

Preference change was also a theme of Richard Bradley's contribution (*Forming a Preference: Incompleteness, Omniscience and Deliberation*). He presented his model for change in preference (or, as he would put it, desirability) and belief, and showed how it could be extended to agents who do not have complete preferences. So doing, he proposed that one can find an equivalent to belief revision contraction operation (which withdraws belief) in the probabilistic framework. Hans Rott's talk (*Defending the Ramsey Test for the Interpretation of Conditionals: A Constructive Approach*) was also concerned with extending concepts from the belief revision literature, though without leaving the AGM paradigm. He showed how using different iterated revision operators borrowed from the belief revision literature, one could, or could not, avoid Gärdenfors famous triviality result regarding the Ramsey Test.

Alexandru Baltag (*Models of Change, Change of Models: Reasoning about the Social Dynamics of "Information"*) argued that Dynamic Epistemic Logic is a fruitful framework for couching traditional epistemological debates, showing how several key concepts in epistemology can be formally expressed in this framework. He also drew links between different sorts of

belief change operators and aggregation functions from social choice theory. Hans van Ditmarsch (*On Knowledge, Knowability and Ability*) also discussed the relationship between Dynamic Epistemic Logic and the philosophical literature, considering the Fitch paradox and different ways of formalising the notion of being able to know something.

The colloquium left participants as any conference of this sort should: with lots to think about. And, who knows, perhaps some changes in attitude!

Brian Hill
IHPST & GREGHEC, HEC, Paris

Ondrej Majer
Philosophy, Academy of Sciences of the Czech Republic, Prague

Calls for Papers

HUMANA.MENTE: Volume 8, Models of Time, deadline 15 November.

SIR KARL POPPER ESSAY PRIZE: British Society for the Philosophy of Science, deadline 31 December.

REASONING FOR CHANGE: Special issue of the journal *Informal Logic*, deadline 10 February.

EXPERIMENTAL PHILOSOPHY: Forthcoming issue of *The Monist*, deadline April 2011.

§4

INTRODUCING ...

In this section we introduce a selection of key terms, texts and authors connected with reasoning. Entries will be collected in a volume *Key Terms in Logic*, to be published by Continuum. If you would like to contribute, please [click here](#) for more information. If you have feedback concerning any of the items printed here, please email thereasoner@kent.ac.uk with your comments.

INDUCTIVE LOGIC

Inductive logic takes various forms, and the simplest characterisation of what these have in common is that they involve inference that is not deductive. Inductive inferences are contingent, that is, the conclusions of inductive arguments do not necessarily follow from the premises; there is no guarantee that true premises lead to true inductive conclusions. Rather, the conclusions are plausible given the premises. Some inductive inferences draw general conclusions from particular cases: from the premise that all emeralds in my experience

have been green, I draw the conclusion that all emeralds are green. Conversely, some draw particular conclusions from general claims: from the premise that all previous raindrops were wet, the conclusion is drawn that the next raindrop will be wet. And some draw particular conclusions from particular cases: from the premise that this ice cube is cold, and this one, and this one, it is concluded that the one over there is also cold. Inference to the best explanation and analogy are commonly used forms of inductive inference.

David Hume argued that inductive inference is not valid, and that there is no reason at all to think the conclusions of such arguments are true; we cannot even say that our inductive conclusions are probably true. Hume claimed that inductive inference is just something that creatures like us find it natural to do, even though it has no logical justification. There is no deductive reason to think that inductive inference is valid—no guarantee that inductive conclusions follow from their premises—and so the only way to justify induction is non-deductively or inductively, but this would be to reason in a circle. Induction depends on the assumption that I have experienced a representative sample of reality, that my limited experience of *F*s is likely to lead to true conclusions concerning all *F*s. But what reason have I to think that this is true? My experience may have been a good guide so far, but to think that it will continue to be so would be to assume that induction is valid, and this begs the question against the Humean sceptic.

Various responses have been offered to the Problem of Induction. Karl Popper accepted that induction is not justified but argued that this is not important because in both everyday reasoning and science we use a form of deductive reasoning instead. In contrast, several attempts have been made to justify induction. It has been argued that inductive inference is by definition rational and thus justified; that if any form of reasoning can identify regularities in nature, induction can, and thus inductive inference is pragmatically justified; and, even though an argument cannot be provided to justify induction, it turns out that inductive inference is reliable—it leads to true conclusions—and thus, again, such reasoning is justified. There is, however, no consensus on whether any of these strategies are successful.

Dan O'Brien
Philosophy, Birmingham

LEWIS, CLARENCE IRVING (1883-1964)

An American philosopher with pragmatist leanings, Lewis criticized the handling of material implication in Russell and Whitehead's *Principia Mathematica* in "Implication and the Logic of Algebra" (*Mind*, 1912). In-

stead, in “The Calculus of Strict Implication” (*Mind*, 1913), he proposed a “strict implication” with ramifications for formal modal languages. His idea of the “pragmatic a priori” opposes necessary truths to the factually contingent rather than to voluntary action or thought.

Kevin S. Decker

Philosophy, Eastern Washington University

§5 EVENTS

OCTOBER

SUM: Second International Conference on Scalable Uncertainty Management, Naples, 1–3 October.

PLACE OF EPISTEMIC AGENTS CONFERENCE: Autonomy and dependence in epistemology, Universidad Carlos III, Madrid, 2–3 October.

MIMS: Workshop on New Directions in Philosophy of Mathematics, Manchester, 4 October.

SETN: 5th Hellenic Conference on Artificial Intelligence, Syros, Greece, 2–4 October.

REASON, ACTIVISM, AND CHANGE: University of Windsor, 3–5 October.

ADVANCES IN CONSTRUCTIVE TOPOLOGY AND LOGICAL FOUNDATIONS: Workshop in honor of the 60th birthday of Giovanni Sambin, Padua, 8–11 October.

FORMAL MODELING IN SOCIAL EPISTEMOLOGY: Tilburg Center for Logic and Philosophy of Science, 9–10 October.

ICAI: The 1st International Conference on Advanced Intelligence, Beijing, 19–22 October.

FoTFS VII: Bringing together Philosophy and Sociology of Science, Foundations of the Formal Sciences VII, Vrije Universiteit Brussel, 21–24 October.

UNCERTAINTY REASONING FOR THE SEMANTIC WEB: 4th International Workshop, in conjunction with the 7th International Semantic Web Conference, Karlsruhe, Germany, 26 October.

MICAI: 7th Mexican International Conference on Artificial Intelligence, Mexico City, 27–31 October.

MDAI: Modeling Decisions for Artificial Intelligence, Barcelona, 30–31 October.

NOVEMBER

PETER LIPTON MEMORIAL CONFERENCE: Department of History and Philosophy of Science, Cambridge, 1 November.

LNAT: Logic Now and Then, The Center for Research in Syntax, Semantics and Phonology (CRISSP), Brussels, 5–7 November.

AUTOMATED SCIENTIFIC DISCOVERY: AAAI Fall Symposium, Arlington, Virginia, 7–9 November.

MWPMW 9: Ninth annual Midwest PhilMath Workshop, 8–9 November.

WPE: Workshop on Philosophy and Engineering, The Royal Academy of Engineering, Carlton House Terrace, London, 10–12 November.

LEBENSWEIT AND LOGIC: The Erlangen school as heir to logical empiricism, Nancy, France, 13–14 November.

NATURE AND STRUCTURE: Philosophy of Physics Graduate Student Conference, SUNY at Buffalo, 15 November.

PROPOSITIONS: ONTOLOGY, SEMANTICS, AND PRAGMATICS: Venice, Italy, 17–19 November.

PHYSICS MEETS BIOLOGY: Perspectives from Philosophy, History, and Science, Royal Society of Edinburgh, 18–20 November.

GAME THEORY: 5th Pan-Pacific Conference in Game Theory, Auckland, 19–21 November.

NEW DIRECTIONS IN EPISTEMOLOGY: International Symposium, Canadian Society for Epistemology, Carleton University, Ottawa, Canada, 21–22 November.

KEAPPA WORKSHOP: Knowledge Exchange: Automated Provers and Proof Assistants, Doha, Qatar, 22 November.

DECEMBER

INFERENCE, CONSEQUENCE, AND MEANING: Sofia, 3–4 December.

ICLP: 24th International Conference on Logic Programming, Udine, Italy, 9–13 December.

CIMCA: International Conference on Computational Intelligence for Modelling, Control and Automation, Vienna, Austria, 10–12 December.

TRENDS IN LOGIC VI: Logic and the foundations of physics: space, time and quanta, Brussels, Belgium, 11–12 December.

ICDM: 8th IEEE International Conference on Data Mining, Pisa, 15–19 December.

PRICAI: Tenth Pacific Rim International Conference on Artificial Intelligence, Hanoi, Vietnam, 15–19 December.

JANUARY 2009

LFCS: Symposium on logical foundations of computer science, Deerfield Beach, Florida, 3–6 January.

SODA: ACM-SIAM Symposium on Discrete Algorithms, New York Marriott Downtown, 4–6 January.

BIOMOLECULAR NETWORKS: from analysis to synthesis, Pacific Symposium on Biocomputing, Fairmont Orchid, The Big Island of Hawaii, 5–9 January.

3RD INDIAN CONFERENCE ON LOGIC AND ITS APPLICATION: The Institute of Mathematical Sciences, Chennai, India, 7–11 January.

GRADUATE CONFERENCE: Second Cambridge Graduate Conference on the Philosophy of Logic and Mathematics, 17–18 January.

VAF: 3th Conference of Dutch Flemish Association for Analytical Philosophy, Tilburg University, the Netherlands, 22–23 January.

BAYESIAN BIostatISTICS: Houston, Texas, 26–28 January.

FEBRUARY

ACM INTERNATIONAL CONFERENCE ON INTELLIGENT USER INTERFACES: Sanibel Island, Florida, 8–11 February.

AIA: IASTED International Conference on Artificial Intelligence and Applications, Innsbruck, Austria, 16–18 February.

INTERONTOLOGY: 2nd Interdisciplinary Ontology Conference Tokyo, Japan, 27 February–1 March.

MARCH

MODELS AND SIMULATIONS 3: Charlottesville, Virginia, 3–5 March.

ADS: Agent-Directed Simulation Symposium, Part of the Spring Simulation Multiconference, San Diego, California, 22–27 March.

CSIE 2009: 2009 World Congress on Computer Science and Information Engineering, Los Angeles/Anaheim, 31 March–2 April.

APRIL

FOUNDATIONS OF MATH: New York University, 3–5 April.

EUROGP: 12th European Conference on Genetic Programming, Tübingen, Germany, 15–17 April.

AISTATS: Twelfth International Conference on Artificial Intelligence and Statistics, Clearwater, Florida, 16–19 April.

MAY

LOGIC OF JOHN DUNS SCOTUS: 44th International Congress on Medieval Studies at Western Michigan University, 7–10 May.

AAMAS: The Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems, Budapest, Hungary, 11–15 May.

PHILOSOPHY AND COGNITIVE SCIENCE: The XIXth edition of the Inter-University Workshop, Zaragoza, 18–19 May.

UR: Uncertain Reasoning, Special Track of FLAIRS, Island, Florida, USA, 19–21 May.

JUNE

ARGUMENT CULTURES: Ontario Society for the Study of Argumentation, Windsor, Canada, 3–6 June.

CNL: Workshop on Controlled Natural Languages, Marettimo Island, Sicily, 8–10 June.

NA-CAP: Networks and Their Philosophical Implications, Indiana University in Bloomington, 14–16 June.

NAFIPS: 28th North American Fuzzy Information Processing Society Annual Conference, University of Cincinnati, Cincinnati, Ohio, 14–17 June.

JULY

METAPHYSICS OF SCIENCE: University of Melbourne, 2–5 July.

ISHPSSB: 2009 meeting of the International Society for the History, Philosophy, and Social Studies of Biology, Emmanuel College, St. Lucia, Brisbane, Australia, 12–16 July.

LOGIC AND HERESY IN THE MIDDLE AGES: Leeds Medieval Congress, 13–16 July.

OCTOBER

EPSA: 2nd Conference of the European Philosophy of Science Association, 21–24 October.

§6

JOBS

ASSISTANT PROFESSOR: Philosophy of science and technology, University of North Texas, deadline 1 October.

THE LUDWIG LACHMANN RESEARCH FELLOWSHIP: Department of Philosophy, Logic and Scientific Method London School of Economics and Political Science, 7 November.

RESEARCH FELLOW: Faculty Of Philosophy, University of Oxford, 7 November.

LECTURER IN PHILOSOPHY OF SCIENCE AND MEDICINE: UCL, Department of Science & Technology Studies, University College London, 17 November.

LECTURESHIP IN PHILOSOPHY: University of Leeds Faculty of Arts, Department of Philosophy, 28 November.

5 RESEARCH POSITIONS: University of Konstanz, 30 November.

ASSISTANT PROFESSOR: Institute of Cognitive Science at Carleton University, 1 December.

ASSISTANT PROFESSOR: Philosophy of social sciences, Université du Québec à Montréal, Montréal, Canada, 1 December.

COURSES AND STUDENTSHIPS

Courses

MSc IN MATHEMATICAL LOGIC AND THE THEORY OF COMPUTATION: Mathematics, University of Manchester.

MA IN REASONING

An interdisciplinary programme at the University of Kent, Canterbury, UK. Core modules on logical, causal, probabilistic, scientific and mathematical reasoning and further modules from Philosophy, Psychology, Computing, Statistics and Law.

MSc IN COGNITIVE & DECISION SCIENCES: Psychology, University College London.

MIND AS MACHINE: Department for Continuing Education, University of Oxford, 1–2 November.

HEALTH IN CONTEXT: A short course in multilevel modelling for public health and health services research, Universidade Nova de Lisboa, Lisbon, Portugal, 10–14 November.

PHILOSOPHY OF PSYCHOLOGY: Bochum / Tilburg, First European Graduate School, Philosophy of Language, Mind and Science, 10–21 November.

SUMMER INSTITUTE ON ARGUMENTATION: University of Windsor, Canada, contact [H.V. Hansen](#) or [C.W. Tindale](#), 25 May – 6 June 2009.

Studentships

4 PHD POSITIONS: Phd Program in Philosophy and Cognitive Sciences at the University Vita-Salute San Raffele, Milan, Italy, 10 October.

PHD SCHOLARSHIPS: Cognitive Science & Philosophy, Macquarie Centre for Cognitive Science (MACCS) at Macquarie University, Sydney, Australia, deadline 17 October.