
THE REASONER

VOLUME 8, NUMBER 3

MARCH 2014

www.thereasoner.org

ISSN 1757-0522

CONTENTS

Editorial

Features

News

What's Hot in ...

Events

Courses and Programmes

Jobs and Studentships

EDITORIAL

It is a great pleasure to present this March issue of *The Reasoner*. I have chosen [Teddy Seidenfeld](#) as my interview partner because of his well-known work on philosophical foundations of statistical reasoning, decision theory, and imprecise probabilities. Teddy is Herbert A. Simon University Professor of Philosophy and Statistics at the Department of Philosophy at Carnegie Mellon University.

The topics of the interview are as follows: After sketching his intellectual autobiography, Teddy points out three interpretations of “imprecise probabilities” and explains important phenomena like “dilation”



which emerge when precise probability theory is generalized into imprecise probability theory. Throughout the interview, Teddy refers the interested reader to relevant literature for digging deeper into these topics.

The photo shows Teddy in company with Isaac Levi (left) and Henry E. Kyburg, Jr. (right). It was taken at the *17th Biennial Meeting of the Philosophy of Science Association* in Vancouver in 2000.

22

22

27

29

31

31

33

34

NIKI PFEIFER

MCMP, LMU Munich

FEATURES

34 Interview with Teddy Seidenfeld

Niki Pfeifer: Teddy, thank you for agreeing to be my interview partner for this issue of *The Reasoner*. Could you tell us something about your intellectual history? How did you become interested in philosophical questions in the foundations of statistics and statistical decision theory?



Teddy Seidenfeld: My interests in statistical inference and statistical decision theory stem from having studied, first with Henry Kyburg Jr. as an undergraduate at the University of Rochester, and then with Isaac Levi as a graduate student at Columbia University. It was an exceptional treat to learn

Henry's original theory of interval valued probability, *Epistemological Probability*, directly from its creator. (For an early version of epistemological probability theory, see chapters 6 and 7 of Kyburg's 1955 Columbia University Ph.D. thesis, *Probability and Induction in the Cambridge School*, which was supervised by Ernest Nagel.) Henry was a wonderful teacher whose friendship I enjoyed for more than 40 years. His personal recommendation was the principal reason why I pursued my graduate studies at his Alma Mater, Columbia University.

At Columbia, more than 40 years ago, I met my good friend Isaac Levi—another graduate of the Columbia Philosophy Department, and another of Ernest Nagel's students—who taught, and continues to teach me important lessons about decision theory. In particular, Isaac's early work using decision theory with epistemic utilities (*Gambling with Truth*, 1967) illustrates why philosophical investigations into problems of induction need to include decision-theoretic analysis—why it will *not* do to try and solve important problems in induction merely with probabilities and conditional probabilities. Isaac urged me to study R.A. Fisher's seminal contributions to statistics, particularly Fisher's enigmatic *Fiducial Probability* and to understand why Fisher thought it solved the problem (Hume's problem) of *Inverse Probability*. I continue to think that that intellectual struggle is priceless for anyone hoping to follow the evolution of statistics in the 20th Century.

For the last 35 years, I've enjoyed regular, sustained exchanges with Jay Kadane (see his *Principles of Uncertainty* (2011)) and Mark Schervish (see his *Theory of Statistics* (1995)), distinguished members of the Carnegie Mellon Statistics Dept. We meet almost every week and keep each other busy and entertained, always trying to improve on our continuing efforts understanding the scope and limits of Bayesian statistics, and statistical decision theory. You will find a sampling of our efforts from the 1990s in our (1999) collection, *Rethinking the Foundations of Statistics*.

NP: You are a founding member of the *Society for Imprecise Probability: Theories and Applications* (SIPTA, established in 2002) and you served as the SIPTA president for four years. What are imprecise probabilities and what can reasoning researchers learn from imprecise probability theories?

TS: As the SIPTA title says, there are multiple interpretations of "Imprecise Probabilities" (IP), and different accounts of how to apply them. Here, let me point out three.

- 1) At the level of an individual decision maker, a modest interpretation of IP theory arises from incomplete elicitation of precise degrees of belief—precise previsions, in de Finetti's jargon. The decision maker may have made coherent, precise assessments for several random variables, but not yet for others. The relation between those precise assessments the decision maker has already made and those yet to be made is captured by de Finetti's *Fundamental Theorem of Probability*—see section 3.10 of his (1974) *Theory of Probability* for a helpful statement. That result provides, for each random variable X that is yet to be assessed, a closed interval—a lower and an upper prevision (a lower and upper probability, when the variable is an indicator function for an event)—that identifies what might be the coherent precise prevision for X that the same decision maker might assess.
- 2) A more ambitious interpretation of "Imprecise Probabilities" arises when IP theory is used as a tool for

(Bayesian) sensitivity analysis: e.g., Robust Bayesian analysis. Canonical Bayesian theory requires a single *prior* probability and single *likelihood function* for modeling the data. Either of these, or both, might be replaced by a set of probabilities, as part of an attempt to study how sensitive "posterior" probability is to variations in either the *prior* or *likelihood*. This idea too has been part of the contemporary statistician's tool-box for many years. A family of IP models that support (both Bayesian and non-Bayesian) robustness analysis are IP theories using *neighborhood-models*, where a single probability distribution is replaced by a set that is focused on that single probability distribution P and where the set is formed using a "neighborhood" of probabilities that are near-by to P . For a general treatment, see P. Huber's (1981) *Robust Statistics*, and for a fine discussion of Bayesian Robustness, see section 4.7 of J. Berger's (1985) *Statistical Decision Theory and Bayesian Analysis* (2nd ed.). Another good source is D.R. Insua's (1990) *Sensitivity Analysis in Multi-objective Decision Making*.

- 3) Still more radical departures from Canonical Bayesian theory can be found in the Imprecise Probability theories due to Kyburg (1974) *Logical Foundations of Statistical Inference*, and Levi (1980) *Enterprise of Knowledge*. The former uses interval valued probabilities based on known frequency information, and it abandons using conditional probabilities to form the basis for updating interval-valued degrees of belief. Unfortunately, Kyburg's Epistemological Probability remains a computationally demanding theory, which has not received the attention it deserves in my opinion. Levi's theory, from his (1980) *Enterprise of Knowledge*, uses convex sets of probabilities and utilities to model unresolved uncertainty and conflicting values, respectively. It relies on Bayesian-styled decision rules for admissibility that do not produce weak orders, which is one aspect that makes it more radically different than the IP theories I've here labeled as type 1 and type 2.

Of course, though more than 20 years old now, P. Walley's (1990) *Statistical Reasoning with Imprecise Probabilities* rightly remains a classic reference on IP theories, as seen primarily from a Bayesian point of view.

NP: You worked with Larry Wasserman on *dilation*. Could you please explain what "dilation" means and why it is important for statistical reasoning?

TS: As early as my graduate student days, while studying with Levi, I became interested in novel aspects of uncertainty that appear when using sets of probabilities but which are absent in the Canonical Bayesian theory, where all probability judgments are determinate and precise. Together with Levi, we observed that, for instance, a classic result about the value of new information (see I.J. Good's (1967) *BJPS* note, *On the principle of Total Evidence*) does not survive the generalization from precise to imprecise probabilities. That is, with IP models of uncertainty, new cost-free information may have negative-value—because the new information is sure to increase uncertainty in a sense that only IP theory can model! Contrary to the Canonical Bayesian result, within IP theory, you'd rather make a terminal decision *in advance* of being able to learn some new (cost-free) evidence. Good responded with his (1974) *BJPS* note, *A little learning can be dangerous*.

In the early 1990s, Larry Wasserman and I, along with then graduate student in Philosophy Tim Herron, made systematic explorations of this phenomenon. Larry and I called it “dilation.” That title is suggestive, using the geometry of IP sets. The anomalous phenomenon occurs when an updated IP set of probabilities for some random variables grows larger (not smaller) by updating each probability in the set using Bayesian conditionalization, based on some new evidence; perhaps, updating using evidence from a planned experiment. And the IP set dilates under conditionalization, no matter which experimental outcomes results. (*Aside:* Dilation is not relegated to Bayesian updating of IP sets. Dilation occurs also using Dempster’s updating rule applied to a Dempster-Shafer Belief Function based, e.g., on an epsilon-contamination IP model.)

As measured by the IP set of probabilities, when dilation is present, the agent’s uncertainty is sure to increase in the light of the new information. So, should the agent plan to avoid such experiments?

Viewed from another perspective, dilation occurs when a collection of Canonical Bayesian agents share some new experimental evidence, but the result is certain to increase (not decrease) their interpersonal disagreements: It is the opposite of the familiar story about the merging of Bayesian posterior probabilities with increasing shared evidence. Of course, dilation brings new consequences for IP theories of experimental design, consequences that are not present in the Canonical Bayesian theory of experimental design.

Dilation is one of several new phenomena about uncertainty that appear when precise probability theory is generalized into IP theory. And what researcher can resist the fun of such surprises?!

NP: What are the most surprising other new phenomena, which might be most interesting for reasoning researchers?

TS: IP theory also provides a novel approach to modeling expert group opinions and cooperative group decisions. For one relevant contrast, consider “pooling rules” for probabilistic opinions. *Linear pooling* replaces a set of probabilistic opinions with one convex combination of those opinions. Because of convexity, the linear pool fails to preserve, e.g., unanimous judgments of probabilistic independence between two events, *A* and *B*. The surface of probabilistic independence between two events, *A* and *B*, is saddle-shaped and has convex (linear) rulings only along sets of probabilities that give *A* (or *B*) constant probability. Using a linearly-pooled probability to replace a set of expert opinions, where each expert judges *A* and *B* independent, typically results in a pooled-probability that makes *A* and *B* probabilistically dependent. That pooled-probability can recommend a decision to pay for an experiment to learn about the event *A* in order then to guide making a decision that depends solely on event *B*. But this is contrary to the unanimous advice of each expert, where each of them judges *A* and *B* probabilistically independent. Their advice is unanimous that the experiment to learn about *A* has no value to guide a decision that depends only on the event *B*. This problem about conditional probability with the *linear pool* is avoided by those *externally-Bayesian* pooling rules, e.g., the *logarithmic pooling rule*, that preserve unanimity among the experts concerning probabilistic independence. However, the *logarithmic pool* can violate unanimity among the experts regarding *unconditional* probability bounds, because it can move the pooled opinion outside the convex hull of the individual expert’s unconditional probabilistic opinions. Both of these defects in pooling can be avoided

by using, instead, an IP set of probabilities to represent the set of expert opinions, rather than trying to find one probability distribution to do all the work.

As an illustration of how the IP approach leads to a different theory of experimental design in the setting of adaptive clinical trials, see *Bayesian Methods and Ethics in a Clinical Trial Design* (1996), J.B. Kadane (ed.), Wiley: New York. Mark, Jay and I illustrate the problem (above) involving the linear pooling rule and independence in section 4 of our (2010) “Coherence choice functions under uncertainty,” *Synthese* 172: 157–176.

NP: Teddy, you are the Herbert A. Simon University Professor of Philosophy and Statistics at CMU. Did you and do you feel obliged to pursue some of Herbert Simon’s research strands?

TS: Acting together as an unstoppable force within CMU, in 1985 Herb, Dana Scott, and Jay Kadane, helped to create CMU’s Philosophy Department. Herb and Dana remained voting members of the Philosophy Dept. throughout their careers at CMU. In Spring 1997, I was honored by my colleagues, and especially by Herb who authorized the award of a “Chair” in Philosophy, in his name. In the spirit of Herb’s longstanding advocacy of models for *bounded rationality*, I used the occasion in 1997 to introduce some early findings about *Measures of Incoherence*, which was then a budding research project, joint with Mark and Jay.

Our idea is to index departures from the ideal of rationality as coherence-of-preference (as in de Finetti’s “Book” argument for coherent pricing) with a normalized rate of sure loss that the less-than-ideal agent’s incoherence allows. The greater the agent’s incoherence the larger her / his (normalized) sure loss. We produced several related papers, summarized in our (2003) “Measures of Incoherence: How not to gamble if you must,” in *Bayesian Statistics 7*.

Last year we returned to the theme of using degrees of incoherence to model less-than-ideal rationality in, “What kind of uncertainty is that?” (2012) *J.Phil.* 109: 516–533. The setting for that work is the old and familiar challenge of how to use personal probability to represent uncertainty about mathematical/logical assertions. Also, we provide some preliminary results about how to reason from within an incoherent position, not knowing just how incoherent YOU are, but nonetheless identifying robust algorithms for improving YOUR rate of incoherence.

As for whether I feel obliged to pursue such research in connection with the Simon Chair, of course, I answer, “No.” Herb was always much too subtle to have it come out any other way!

NP: Thank you very much for this informative and interesting interview!

Deduction, Novelty, Popper

Danny Frederick denies what I shall call ‘the containment principle’, namely, that “the conclusion of a valid deductive argument is contained in its premises and says nothing new” (*The Reasoner*, 5.4, pp. 56-57).

His denial pursues an older controversy. We are taught that deductive validity is a property of the form of an argument. In a valid deductive form, the propositions in the conclusion are contained in the premises. Thus a validly derived conclusion appears merely to repeat what is in the premises. If so, how can deduction generate new information?

Ayer's early answer (1936, *Language, Truth and Logic*, 2001, Penguin, pp. 73–75) is that deduction doesn't generate new *objective* information, rather it brings into one's subjective awareness information which was always present but inaccessible.

Hintikka's response conceives new information as an increase in degrees of belief with a corresponding reduction of uncertainty (Hintikka, 1970, 'Information, Deduction and the A Priori', *Nous*, 4, 2, pp. 135–152). On this account since 'degrees of belief' and 'uncertainty' are conceived objectively, deduction is thought to add to the stock of objective information.

Frederick's denial of the containment principle offers a third solution. For support he co-opts Popper, who, Frederick maintains, refuted the principle in his *Unended Quest* (1976, Fontana, pp. 25–28). But Popper neither refuted the containment principle, nor intended to; rather, he is committed to it. The refutation is supposed to be found in the following counter-example (*The Reasoner* 5.4, p. 57):

"Let N stand for Newton's theory of gravitation and let E stand for Einstein's theory of gravitation. Since N is incompatible with E , the following argument is deductively valid:

N , therefore, not- E .

But the conclusion of this argument would certainly have said something new in Newton's time. Newton could not foresee Einstein's theory; and none of his contemporaries could have arrived at a statement of the negation of Einstein's theory simply by unfurling the implicitly known content of Newton's theory."

Notice that

N , therefore, not- E

is not a valid argument. It is made valid by adding a premise asserting the incompatibility of N with E . Einstein supplied E , yielding the conditional,

If N then not- E .

The argument is now valid, but is so because the conclusion not- E is now contained in the premises. Popper's argument is therefore not a counter-example to the containment principle, for it is not an instance of a valid deductive form in which the conclusion is not contained in the premises.

Is there an interpretation which makes Frederick's denial plausible? The answer lies in isolating then disambiguating two distinct propositions which Frederick's complex statement fudges. He denies two propositions:

P1: The conclusion of a valid deductive argument is contained in its premises.

P2: The conclusion of a valid deduction says nothing new.

In P2, 'says nothing new' is ambiguous between:

- (a) expresses no proposition not expressed before *this time*,
and
- (b) expresses no proposition not expressed before *in the deduction*.

Understood as (a), we can agree with Frederick that since Einstein's theory was unknown in Newton's time, its denial *then* would say something new. Frederick appears to believe this sufficient to refute the containment principle, but this belief is mistaken. Understood as (b), P2 is true: having occurred in the conditional premise, in the conclusion not- E says nothing new *in the deduction*. This suggests where Frederick may err: the containment principle states a relation between propositions in a valid deduction; it says nothing about the awareness of particular reasoners, nor about what is new beyond the domain of the deduction.

Had Popper sought to refute the containment principle we would expect him to hold that not- E is not contained in N . Contrariwise, he states that not- E "belongs to the *logical content* of N " (Popper, p. 27, italics original). Here, 'belongs to' and 'content' reveal that contra Frederick, Popper holds that not- E is contained in N . He identifies what a statement entails with what it contains (Popper, p. 26) and since in a valid deduction the conclusion is entailed by, and hence contained in, the premises, Popper is committed to the containment principle.

There is also cause to doubt Frederick's claim that Newton and his contemporaries "could not have arrived at a statement of the negation of Einstein's theory simply by unfurling the implicitly known content of Newton's theory." This formulation conflates the content of Newton's understanding and the content of his theory. True, Newton couldn't have 'implicitly known' that Einstein would theorise gravity as an effect of space-time curvature, but by deduction Newton could have concluded that N entails $N1$,

$N1$: massless objects are unaffected by gravity.

Einstein's theory E entails $E1$,

$E1$: massless objects are affected by gravity.

Newton needn't know anything about E to know that $N1$ entails not- $E1$. Since $E1$ is entailed by E , not- $E1$ entails not- E ; hence $N1$ also entails not- E . And since $N1$ is entailed by N , N entails not- E . To negate a statement is to entail that it is false, therefore, although ignorant of E , Newton did in fact arrive at statement N which negates Einstein's theory and any other theory or statement with which N is inconsistent.

How then does deduction generate new information? The special merit of deduction is its truth-preserving character: the truth of the premises carries into the conclusion. This occurs only because at least one of the premises carries into the conclusion. Frederick's solution fails because it deprives validity of this character. A better answer goes like this: observation, ingenuity, and new theories supply new premises (Einstein's theory and subsequent observations of the apparent gravitational influence on light are paradigmatic examples). Second, by combining new or long strings of premises, deduction unveils formerly obscure relationships. Third, as Frederick observes, Popper taught that the content of a theory is infinite; deduction reveals elements of that content. Each of these mechanisms reduces our uncertainty about our world and in this sense, Hintikka's sense, yields new information. However, if 'uncertainty' and 'degrees of belief' are subjective mental attitudes not objective probabilities, then even on Hintikka's account, new objective information emerges only from new premises, not from deduction.

DAVID MCBRIDE
Philosophy, Open University

Against Phenomenal Conservatism: a Reply to Moretti

According to Phenomenal Conservatism (Huemer, 2007, "Compassionate Phenomenal Conservatism," *Philosophy and Phenomenological Research* 74, 30–55):

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

In (Mizrahi, 2013, "Against Phenomenal Conservatism," *The Reasoner*, 7(10), 117–118), I advance a *reductio* against (PC) which shows that appealing to seemings is an untrustworthy method of fixing belief (MFB). The following is a key premise in my *reductio* against (PC):

- (4) If an MFB provides at least some degree of justification for contradictory beliefs, then it is untrustworthy.

An anonymous reviewer for *The Reasoner* has offered the following putative counterexample to (4):

Suppose that you have an urn and know that it contains a Red, a Blue and a Yellow ball. Alice takes out one ball (you cannot see or otherwise know its colour). When Alice truthfully tells you that the ball is not Yellow, then this information gives you reason to believe that it is Red, and also reason to believe that it is Blue. The beliefs 'the ball is Red' and 'the ball is Blue' are clearly contradictory. But this does not make the source of justification untrustworthy.

Moretti (2013, "Mizrahi's Argument against Phenomenal Conservatism," *The Reasoner*, 7(12), 137–139) shows that this putative counterexample doesn't challenge (4) because, although incompatible, (r) 'it is red' is not a logical negation of (b) 'it is blue'. Moreover, (e) 'it isn't yellow' doesn't give one some justification for both (r) and $\neg(r)$ [or for both (b) and $\neg(b)$] because, after one learns (e), one's degree of confidence in $\neg(r)$ drops from 2/3 to 1/2. If (e) increases one's degree of confidence in (r), it must decrease one's degree of confidence in $\neg(r)$.

However, Moretti doesn't think that my *reductio* against (PC) is free of problems. He argues that (4) is implausible because it implies that perception and testimony are untrustworthy MFBs. Since he thinks that perception and testimony are trustworthy, Moretti (2013, 138) amends (4) as follows:

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

With (4*), Moretti argues, my *reductio* against (PC) becomes invalid, since S_1 's seeming is not the same evidence as S_2 's seeming.

I think that Moretti makes a good point that can be accommodated by my *reductio* against (PC). That is, S_1 's seeming that p and S_2 's seeming that $\neg p$ are the same *type of evidence*, namely, seemings, but with distinct contents, i.e., $\langle p \rangle$ and $\langle \neg p \rangle$. That makes S_1 's seeming that p and S_2 's seeming that $\neg p$ different *pieces of evidence* of the same type. To see why, consider the litmus test case. Although the litmus paper turning

red and the litmus paper staying blue are distinct *pieces of evidence*, they are the same *type of evidence* because they are produced by the same method, namely, the litmus test. Similarly, Jackson's seeming and Dennett's seeming are distinct pieces of evidence of the same type, namely, seemings.

Now, if an MFB produces *distinct pieces of evidence of the same type*, which provide some degree of justification for contradictory beliefs, then it's untrustworthy. This is because distinct pieces of evidence of the same type, which are produced by the same MFB, are neither rebutting nor undercutting defeaters (Pollock, 1987, "Defeasible Reasoning," *Cognitive Science*, 11, 481–518). Instead, *they undermine our confidence in the method itself*. To see why, consider an MFB that Moretti deems trustworthy, namely, testimony. Suppose two eyewitnesses, E_1 and E_2 , give contradictory testimonies. In this case, E_1 's testimony that p is a *prima facie* reason for p and E_2 's testimony that $\neg p$ is a *prima facie* reason for $\neg p$. However, E_1 's testimony that p is a rebutting defeater for $\neg p$ and E_2 's testimony that $\neg p$ is a rebutting defeater for p . We also have a reason for denying that ' E_1 says that p ' is a reason for p and that ' E_2 says that $\neg p$ ' is a reason for $\neg p$, given that it's not the case that p wouldn't be true unless E_1 said so or that $\neg p$ wouldn't be true unless E_2 said so (Pollock 1987, 485). More importantly, however, we now have two pieces of evidence (i.e., E_1 's testimony that p and E_2 's testimony that $\neg p$) of the same type (i.e., testimony) that provide some degree of justification for contradictory beliefs. For this reason, we wouldn't trust the testimonies of E_1 and E_2 in this case. In other words, appealing to eyewitness testimony is an untrustworthy MFB in this case.

To address Moretti's criticism, then, my *reductio* against (PC) can be revised as follows:

1. (PC) [Assumption for *reductio*]
2. It seems to S_1 that p and it seems to S_2 that $\neg p$, independently of each other. [Premise]
3. \therefore In the absence of defeaters, S_1 has some degree of justification for believing p and S_2 has some degree of justification for believing $\neg p$. [From (1) & (2)]
4. If an MFB produces distinct pieces of evidence of the same type that provide some degree of justification for contradictory beliefs, then it's untrustworthy. [Premise]
5. \therefore Appealing to seemings produces distinct pieces of evidence (a seeming that p and a seeming that $\neg p$) of the same type (seemings) that provide some degree of justification for contradictory beliefs. [From (3)]
6. \therefore Appealing to seemings is an untrustworthy MFB. [From (4) & (5)]

Just as we wouldn't trust the testimonies of E_1 and E_2 , since we have two pieces of evidence (i.e., E_1 's testimony that p and E_2 's testimony that $\neg p$) of the same type (i.e., testimony) that provide some degree of justification for contradictory beliefs, we wouldn't trust the seemings of S_1 and S_2 , since we have two pieces of evidence (i.e., S_1 's seeming that p and S_2 's seeming that $\neg p$) of the same type (i.e., seemings) that provide some degree of justification for contradictory beliefs. In other words, appealing to seemings is an untrustworthy MFB.

MOTI MIZRAHI

Philosophy, St John's University

The Social Dynamics of Information Change, 2–4 December

This workshop was the second major event associated with Sonja Smets' ERC funded project "The Logical Structure of Correlated Information Change". The aim of the workshop was to create a platform for researchers in many areas to present their work and exchange ideas on the social dynamics of information change. The workshop consisted of three days worth of talks, with each day covering a cluster of topics.

Day one was broadly centered on social networks and epistemic logics.

Some of the talks were methodological in flavour: Jeroen Bruggeman discussed the sociological issues surrounding gossip within social networks, and their links to network topology and social cohesion; Vincent Hendricks proposed a unifying framework for a variety of socio-informational phenomena, and then discussed the pivotal role of public signals therein and the complications introduced by their ambiguity; and Christian List presented a new framework that allows a unified formal representation of a large swathe of moral theories.

Several talks focused on epistemic logics for modeling social effects: Rasmus K. Rendsvig argued that threshold models do not fully capture information sharing in social networks, and then presented an alternative model that, he argued, does; Alexandru Baltag presented a number of modal formalisms that capture various dynamic-informational features of social networks; Hans van Ditmarsch presented an epistemic logic wherein facts about agents' awareness of propositional variables are expressible, and then compared various epistemic notions that arise in the presence of awareness; and finally, Aybüke Özgün presented a new topological semantics for Stalnaker style belief logics.

The remainder of the day's talks focused on relevant computational considerations: Eric Olsson defended a Bayesian theory of Trust in social networks, and used this in conjunction with the simulation environment LAPUTA to formulate a hypothesis about the value of overconfidence in social networks; Thomas Bolander presented a proof of the undecidability of the multi-agent epistemic planning framework that Dynamic Epistemic Logic provides; Jan van Eijck presented an extension of Dynamic Epistemic Logic intended for model checking cryptographic protocols, and showed how it can be used for model checking the Diffie-Helman key exchange and similar protocols.

Day two was centred on the higher-order reasoning involved in social interactions.

The first block of the day focused on linking abstract logical theories and empirical work on the theory of Mind: Rineke Verbrugge discussed human aptitude in higher-order reasoning tasks and environmental factors leading to the emergence of higher-order reasoning; Jakub Szymanik argued that in classifying the computational feasibility of socio-epistemic problems, one should take the perspective of the bounded agents involved, rather than the modeler's perspective.

Quite a few talks were devoted to the relevant game-theoretical accounts: Amanda Friedenberg presented a characterization of the behavioral implications of rationality plus any finite order of strategic reasoning; Paolo Galeazzi compared type spaces and Kripke structures as models for game theory,

and explored their relationships; Andrés Perea presented a solution concept for dynamic games whereunder reasoning is temporally local, discussed its relationships to classical solution concepts, and provided a procedure for computing strategies in accordance with the new concept; and in the final lecture of the day, Johan van Benthem discussed reasoning about strategies, proposed several logics that analyze elementary game theoretic reasoning, and identified some challenges for new research in this area.

Day three started with some complementary talks on the contribution of logic to social network theory.

Jeremy Seligman presented various dynamic hybrid logics for representing social network structures, and their dynamics. Jens U. Hansen presented a logical framework for representing the flow of information through a social network in an way intermediate between local and global perspectives, which Zoé Christoff followed-up by applying this framework to cases of pluralistic ignorance.

The afternoon session was devoted to the formal learning theory accounts of reasoning and interaction: Kevin Kelly presented a modal semantics for inductive knowledge whereunder agents are computable learners who modify their beliefs over time based upon new data and the prior beliefs of other agents; Nina Gierasimczuk discussed the links between logics of the social dynamics of information and computational learning theory, with a view to shedding new light on certain research agendas in logic as the study of information processing in the context of learning, which Dick de Jongh followed up by explaining how finite identifiability in learning theory can be used to study the problem of conclusive knowledge update; Thomas Zeugmann presented results relevant for detecting social dynamics of information change, particularly when applied to internet phenomena, and also results on learning algorithms based upon ultrametrics; Éric Martin presented a general logical framework for informational change which refines the model-theoretic and proof-theoretic notions of compactness and inference, and used this to inspire a generalization of modal logic; and the final talk of the workshop was given by John Case who presented results about reactive learners that learn from one another in an interactive way.

As the reader may have already noted, the workshop talks looked at both well-established work and possible new connections between logic, belief revision theory, learning theory, game theory, and social science. The hope is that, in bringing together ideas from these different areas, new theoretical work shall arise, which shall entrain new modeling techniques and a better understanding of puzzling social-informational phenomena such as pluralistic ignorance, the bandwagoning effect, group polarization, and more: fingers crossed.

CALLUM SIDA-MURRAY

Institute for Logic, Language and Computation, Amsterdam

Conditional Thinking, 14–15 January

The ERC-funded Nature of Representation Group hosted a conference on [Conditional Thinking](#) at the University of Leeds. Speakers included A.M. Ahmed, Richard Bradley, Dorothy Edgington, Daniel Elstein, Hlynur Orri Stefánsson, and J. Robert G. Williams.

Dorothy Edgington's talk, "The Suppositional Theory of Conditionals and the Frege-Geach Problem," began the confer-

ence with a survey of ways in which different accounts of conditionals handle embeddings. Like other forms of non-truth-conditional views (for example, in the analysis of normative language), non-truth-conditional accounts of conditionals face a challenge for delivering a systematic and predictive account of how conditionals are embedded in other apparently truth-evaluable constructions and how to predict when they'll be judged true and false. Edgington argued that truth-conditional accounts of modals face the second problem equally. On her view, Bradley's two-dimensional account of conditionals handles these challenges best and is amenable to an anti-realist interpretation.

Arif Ahmed's talk, "Modality and Prudence," focused on a central conference theme: the question of what role counterfactual conditional thinking plays in our mental economy. Ahmed argued for a fundamentally pragmatic role for counterfactual thinking. In assessing our own choice behavior, we entertain counterfactual thoughts like: "If I had taken the second box, I would have another \$1000." These thoughts generate regret or its positive counterpart; and these affective responses reinforce more rational choice behavior in future decision-making.

Daniel Elstein and J.R.G. Williams's talk, "Decisions and Suppositions," began by distinguishing two forms of supposition. "A-supposition" involves treating a proposition as if it is new evidence for the (imagined) reconfiguration of beliefs; "C-supposition" involves an imagined reconfiguration of the world. Elstein and Williams argued that the central role for C-supposition is for reasoning about action. The norms that govern this form of supposition are not norms based on probabilities of counterfactuals, but rather on conditional chances.

Richard Bradley and Orri Stefánsson's talk, "Counterfactuals, Decisions and Chances," presented a multidimensional framework for counterfactuals. Counterfactual conditionals are made true by "counterfactuals"; elementary possibilities are ordered tuples of worlds: a world of evaluation and a sequence of "counter-actual" worlds. Bradley and Stefánsson argued that this framework (first introduced in Bradley, "Multidimensional Semantics for Conditionals," *Philosophical Review* 121 (2012) and "Conditionals and Supposition-Based Reasoning," *Topoi* 30 (2011)) has interesting applications in decision theory. Their claim: what could have happened sometimes affects the desirability of what actually happens. This is used to provide an unorthodox explanation of the Allais and Diamond paradoxes.

JENNIFER CARR

Philosophy, University of Leeds

Causality in the Biological Sciences, 17 January

In January this year C. Kenneth Waters (Center for Philosophy of Science, University of Minnesota) came to Cologne as a visiting professor of our DFG-research group "Causation and Explanation". One highlight of the fruitful interactions between Ken and our group was the workshop on "Causality in the Biological Sciences" that took place at the University of Cologne on January 17th, 2014. Besides Ken Waters, the invited speakers were Lorenzo Casini (LMU München/Université de Genève), Kolja Ehrenstein (Universität zu Köln), Lena Kästner (Ruhr-Universität Bochum), and Raphael Scholl (Universität Bern).

The aim of this workshop was to bring together researchers who work on various aspects of causal reasoning in the life sci-

ences. Some of the talks focused more on general philosophical issues, such as the problem of omissions / absences, the locality of causation (Ehrenstein), and the question of how to model mechanisms by means of Bayesian networks (Casini). Other speakers discussed how causal reasoning and explaining works in specific scientific fields, such as classical and contemporary genetics (Waters), neuroscience (Kästner), and biochemistry (Scholl). The workshop also brought together work on the history of causal reasoning in the biological sciences (Scholl, Waters) and more systematically-oriented work on causality (all).

Ken Waters opened up the workshop with a talk about "Causes that Matter in Scientific Practice". The central question of his talk was: What makes DNA so valuable for biological research? Ken argued that in many contexts DNA is important because it is the *actual difference maker* of nucleotide sequence differences in RNA and polypeptides. However, he also pointed out that the process from DNA to a functioning protein is complex, which is why there are various limitations of the explanatory significance of DNA as an actual difference maker. Furthermore, Ken used an example of investigative practice from genetics to show that the value of DNA consists not so much in its explanatory significance, but rather in its utility as a means for manipulation. In these contexts DNA is, in Ken's terms, a *practical potential difference maker*.

Raphael Scholl presented a project on "Discovery from a Causal Point of View: Oxidative Phosphorylation", which he is pursuing together with Karin Nickelsen (LMU München). He used a historical case study (namely Peter Mitchell's discovery of the chemiosmotic theory of oxidative phosphorylation) to point out how a handful of simple causal heuristics suffice to explain the genesis of even very original, Nobel-prize-worthy hypotheses. Another major result of Raphael's analysis was that only such causal structures were investigated (or mentioned) for which the underlying mechanisms were known, but that the coarse-grained causal level developed much more continuously than the fine-grained mechanistic level.

In the afternoon Lorenzo Casini presented recent developments of his project "How to Model Mechanisms: In Defense of Recursive Bayesian Networks". He works on this project together with Jon Williamson (University of Kent). The central goal of Lorenzo's project is to show how Bayesian nets can be used to model biological mechanisms (such as the mechanism of apoptosis) and how putative problems can be solved. One challenge that formal models of mechanisms encounter is how to account for causal reasoning across levels. Lorenzo's solution is to use recursive Bayesian nets (RBNs) to represent hierarchies of mechanisms. In his talk he responded to objections and revealed the advantages that RBNs have over other accounts.

Lena Kästner's talk on "Materials & Methods" addressed the question of how scientists develop causal explanations. Her main aim was to point out the limited role that interventions (of the Woodward style) play in finding causal explanations. On the basis of various examples Lena showed that non-intervention strategies like mere interactions, pseudo-interventions, and data analyses are crucial to causal inference in neuroscience.

The last talk of this workshop was given by Kolja Ehrenstein. He discussed "How to and not to break locality". Kolja examined whether and which understanding of locality may help to qualify absences, such as the absence of lactose (which is supposed to cause the absence of β -galactosidase in the *lac*

operon), as causes (or as effects). He argued that locality, understood as spatiotemporally continuous sequences of causal intermediates, fails. Instead, he defended the view that negative causes should be understood as disconnections of dispositional overlap and negative effects as inhibitions of the manifestation of a disposition.

I would like to thank all the speakers and participants of this workshop, in particular Ken Waters, for their stimulating contributions. Information about future events of our research group can be found on our [website](#).

MARIE I. KAISER

Philosophy, Universität zu Köln

Calls for Papers

RTCiSS: Rethinking Theory Construction in Social Science, London School of Economics, 11 March.

CAUSAL DISCOVERY AND INFERENCE: special issue of *ACM Transactions on Intelligent Systems and Technology*, deadline 14 March 2014.

THE HARD PROBLEM OF CONSCIOUSNESS: special issue of *Topoi*, deadline 28 March.

PRESUPPOSITIONS: special issue of *Topoi*, deadline 15 May 2014.

VIRTUES & ARGUMENTS: special issue of *Topoi*, deadline 1 September 2014.

WHAT'S HOT IN . . .

Logic and Rational Interaction

Game theory is a tool for modeling strategic reasoning. Its use is prescriptive. Given an interactive situation, game theory analyzes what an agent should do to maximize his outcome. However, here is also a second, less known application of game theoretic modeling: reconstruction.

We sometimes find ourselves observing some agents, stock brokers, football players or diplomats, that we deem to be competent in what they do. Being new to the field, we might not know what exactly it takes to be successful; we don't know enough about the strategic situation the agents are engaged in. We do observe behavior, but we don't know the game. Assuming the agents are rational, we can then try to reconstruct the situation from their behavior. A classical example of such reconstructive use is Robert Putnam's seminal 1988 paper on [two layer games](#). During the seventies, the western world underwent a major economic downturn following the oil crisis. In the aftermath of this crisis several of the G7 group of nations tried to coordinate on joint action at several economic summits. Political scientists attempted to explain the coordinative outcomes of these conferences by considering the strength of the individual actors, both on economic and political grounds. However, single negotiation models had failed at explaining the outcomes of the 1978 Bonn summit. Putnam's seminal idea was to devise the theory of two level games. Having coalition partners and national parliaments sitting on their backs, the individual negotiation partners were not involved in one negotiation process, but in two simultaneous processes, one on the international level and a second one in their home game. This led to the seemingly paradoxical outcome that actors with a strong domestic standing would need to give in more than their counterparts with weaker positions, simply because it was

commonly known that these agents were able to defend more compromise outcomes against their domestic stakeholders.

A similar reconstructive approach underlies the idea of signaling games. Nature is full of signals between species. Baby birds begging for worms, flies looking like wasps to scare predators, peacocks growing big tails to appear as attractive mating partners. Assuming that the species playing the better strategy is more likely to reproduce, we can assume that nature moves slowly towards a Nash equilibrium in the signaling game.

The big question then is: How did signals emerge the way they did? If growing big tails increases the chances of mating for a peacock, then why don't all peacocks do so? The traditional game theorist's answer is: Because it's costly. Having a big tail makes it more difficult to escape predators. And certainly, being eaten isn't helpful in passing on your genes. Weak peacocks don't signal strength, because they cannot afford it. Now, there are cases where sending the signal is very cheap. In these cases, why don't the weak agents start mimicking the strong agents' signal? Indeed baby birds can be observed begging for food, even if they are most likely not needy. In this case we could even expect the signal to vanish completely, since it does not carry any information any more.

In a recent paper, [Kevin Zollman](#), [Carl Bergstrom](#) and [Simon Huttegger](#) show that signals can emerge as long as sending a signal is cheaper for strong agents than for weak agents. Basically, the receiver, the bird parents, can decrease the potential payoff of a signal by not replying to every incoming signal, but randomizing their answering behavior. For instance, mexican blue footed boobies have been observed not replying to begging immediately, but waiting for up to 20 minutes before giving in to their offspring. As it turns out, the Nash equilibrium, that is a state in which both, sender and receiver act optimally given the opponent's behavior, is such that strong agents always signal truthfully while weak agents mix between signaling and not signaling.

Now, if we are presented with a signaling situation from nature, what is the right reconstruction? The original, high cost analysis, or the novel analysis where weak agents face low cost for the signal. If we observe weak agents mixing strategies, is it because this is a close to optimal move in a low cost game, or has the evolutionary process simply not yet converged towards an equilibrium state? The authors of the paper give an answer to this question in terms of evolutionary data: in the original, high cost analysis, the population should display a uniform dynamics towards the equilibrium, while in the low cost case the strategies will most likely display a cycling behavior around the equilibrium value.

LORIWEB is always happy to publish information on topics relevant to the area of Logic and Rational Interaction—including announcements about new publications and recent or upcoming events. Please submit such news items to Carlo Proietti, our web manager via



DOMINIK KLEIN
TiLPS, Tilburg University

Philosophy of Cognitive Science

The *Epistemic Innocence Project* investigates the potential epistemic benefits of imperfect cognitions, that is, cognitions that are factually inaccurate in some key respect. Imperfect cognitions may include delusional beliefs, distorted memories, and confabulatory explanations in the clinical and non-clinical population. It is currently funded by a 12-month AHRC Fellowship awarded to Professor Lisa Bortolotti at the University of Birmingham, and it also features Ema Sullivan-Bissett as a research fellow. Research questions include the following:



- (a) In what circumstances do delusional beliefs, distorted memories, and confabulatory explanations contribute to the acquisition and retention of true beliefs?
- (b) Do delusional beliefs, distorted memories, and confabulatory explanations have genuinely epistemic benefits?
- (c) Are people epistemically blameworthy for having imperfect cognitions?
- (d) What are the consequences of acknowledging that delusional beliefs, distorted memories, and confabulatory explanations can be epistemically advantageous?

One of the main objectives of the project is to develop an account of epistemic innocence for imperfect cognitions. Ideally, we would have cognitions that satisfy norms of truth and accuracy and that are supported by, and responsive to, the evidence available to us, as well as fostering the acquisition and retention of true beliefs. But we have limited cognitive capacities, and imperfect cognitions that are false or inaccurate and badly supported by, or irresponsive to, the evidence are a common occurrence. The project explores the possibility that some of these imperfect cognitions are epistemically innocent, where the notion of innocence captures the fact that for a given agent at a given time it is epistemically advantageous to have such cognitions, even if they fall short of key epistemic norms. The central idea is that a cognition is innocent if the following two conditions are met:

1. **Epistemic Benefit.** The cognition delivers some significant epistemic benefit to a given subject at a given time, that is, it contributes to the acquisition and retention of true beliefs.
2. **No Relevant Alternatives.** Alternative cognitions that would deliver the same epistemic benefit are unavailable to that subject at that time.

Different notions and degrees of unavailability apply. In general terms, there may be no genuine alternative to an imperfect cognition, because information that would lend support to a different, more accurate, cognition is opaque to introspection, not

open to investigation, irretrievable, or blocked for motivational reasons; or the alternative cognition could be strictly speaking available, but it would not carry the same epistemic benefit as the imperfect cognition (e.g., it would not offer a plausible explanation of the subject's experience or it would not support the subject's sense of self to the same extent).

The themes of the project will be explored in a two-day workshop at the University of Birmingham, entitled "Costs and Benefits of Imperfect Cognitions", to be held on 8th and 9th May 2014. The workshop will be one of the means by which the Epistemic Innocence project interim results are disseminated, and will promote exchange between philosophers and psychologists on the potential pragmatic and epistemic benefits and costs of beliefs, memories, implicit biases, and explanations. Speakers include the project team (Lisa Bortolotti and Ema Sullivan-Bissett at the University of Birmingham), and then Katerina Fotopoulou (University College London), Martin Conway (City University London), Ryan McKay (Royal Holloway) and Maarten Boudry (University of Ghent), Miranda Fricker (University of Sheffield), Jules Holroyd (University of Nottingham), Petter Johansson and Lars Hall (University of Lund).

More information about the project and the workshop can be found on our [website](#), on [twitter](#), on [facebook](#), and on our [blog](#).

LISA BORTOLOTTI
Philosophy, University of Birmingham

Uncertain Reasoning

Galileo Galilei was born on 15th February 1564, in the heart of the historic town centre of Pisa. Celebrations of this notable birthday have taken place in Pisa, Florence and many other locations throughout Italy. The Scuola Normale Superiore was to host, a few hundred meters away from Galileo's birthplace, the presentation of a new commemorative stamp. But uncertainty took over and the celebration had to be canceled the day before. For the Istituto Poligrafico della Zecca Italiana (the Italian Mint) did not manage to get the stamp out on time. Generic "technical problems", they said. Not exactly flattery to one of the main contributors to the birth of modern science.



Interestingly enough (for the purposes of this column!) Galileo also deserves to be celebrated as an uncertain reasoner. In *Considerazione sopra il Giuoco dei Dadi*, an unpublished note dated between 1613–23, he calculated correctly the chances of throwing three dice. The problem was apparently posed to Galileo by a group of Florentine nobles, who had observed a regularity they had no means to explain. Whilst 9 and 10 are obtained from the same number of 3-partitions each—namely, for 10: (631), (622), (541), (352), (442), (433)—experience suggested to (presumably heavy) gamblers that 10 was more likely than 9. By calculating correctly the permutations of the partitions, Galileo showed that 10 appears in 27 out of the 216 possible outcomes, whereas 9 only 25. His key observation to get the result was to acknowledge the independence of the three dice:

since a die has six faces, and when thrown it can equally well fall on any one of these, only six throws can be made with it, each different from all the others. But if together with the first die we throw a second, which also has six faces, we can make 36 throws each different from all the others, since each face of the first die can be combined with each of the second.

A very detailed reference to find out more about Galileo's note is David, F.N. (1962: *Games Gods and Gambling*, Hafner Publishing Company, New York) which also contains (Appendix 2) an English translation of the *Considerazione*. A more recent, and compact, appraisal is provided by Hald, A. (1990: *History of Probability and Statistics and Their Applications before 1750*, John Wiley & Sons).

Historians don't seem to have reached a consensus on whether Galileo was influenced by Tartaglia and Cardano, or if he rather worked out the solution by himself. Yet, as David (1962) points out in connection to the above quotation, Galileo doesn't justify the identification of the probability of an outcome with the ratio of the "favourable cases" to the "equipossible" ones. This suggests that by the 1620s the so-called *classical definition* of probability was widely accepted. If this is correct, what Galileo thought of as his contribution is the above-quoted analysis of independence. This notion would be the fulcrum of subsequent work by Huygens, paving the way to the seminal *Ars conjectandi* by Jacob Bernoulli.



HYKEL HOSNI

Marie Curie Fellow,
CPNSS, London School of Economics

EVENTS

MARCH

- PLUK:** Philosophy of Language in the UK, University of Leeds, 7–8 March.
- WBEM:** Workshop on Beauty and Explanation in Mathematics, Umeå University, Sweden, 11–12 March.
- EMLP:** Empirical Methods of Linguistics in Philosophy, Technische Universität Dortmund, 13–14 March.
- JUSTIFICATION:** Towards an Epistemology of Understanding: Re-thinking Justification, Bern, 21–22 March.
- CELL:** 2nd Logic and Language Conference, London, 21–22 March.
- PREFER:** Respecting Context-dependent Preferences, Umeå University, 21–22 March.
- CORCON:** Correctness by Construction, Genoa, Italy, 24–27 March.
- MIND:** 9th Mind Network Meeting, Oxford, 26 March.
- NNPS:** Nordic Network for Philosophy of Science, Lund University, Sweden, 27–28 March.

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

PBUK: Philosophy of Biology in the UK, Christ's College, University of Cambridge, 31 March–1 May.

APRIL

- NAG:** Norms, Actions, Games, London, 1–2 April.
- AISB:** 7th AISB Symposium on Computing and Philosophy: Is computation observer-relative?, Goldsmiths, London, 1–4 April.
- HAPOP:** History and Philosophy of Programming, Goldsmiths, University of London, 1–4 April.
- D& MC:** Deductive and Mathematical Cognition, Bristol, 7–8 April.
- EBL:** 17th Brazilian Logic Conference, Petrópolis, Brazil, 7–11 April.
- PSX4:** Philosophy of Scientific Experimentation 4, Pittsburgh, PA USA, 11–12 April.
- PHILOSTEM:** 6th Midwest Workshop in the Philosophy of Science, Technology, Engineering, and Mathematics, Fort Wayne, Indiana, 11–12 April.
- MATHEMATICAL DEPTH:** University of California, Irvine, 11–12 April.
- L & MS:** workshop on Logical and Modal Space, New York, 11–13 April.
- TAMC:** 11th Annual Conference on Theory and Applications of Models of Computation, Anna University, Chennai, India, 11–13 April.
- LMS:** London Mathematical Society Lectures by Jouko Väänänen, London, 14–17 April.
- PhML:** Philosophy, Mathematics, Linguistics: Aspects of Interaction, St. Petersburg, Russia, 21–25 April.
- PhDs in LOGIC:** Utrecht, The Netherlands, 24–25 April.
- MAICS:** 25th Modern Artificial Intelligence and Cognitive Science Conference, Gonzaga University, Spokane, WA, USA, 26–27 April.
- UConnLogic:** Abstractionism / Neologicism, University of Connecticut, 26–27 April.

UK-CIM: UK Causal Inference Meeting (UK-CIM): Causal Inference in Health and Social Sciences, University of Cambridge, Cambridge, 28–29 April.

GIRLS: 3rd Conference on Games, Interaction, Reasoning, Learning & Semantics: Evolution and Cooperation, Lund, 28–30 April.

RSC: Research Students' Conference in Probability and Statistics, Nottingham, 28 April–1 May.

UK-CIM: workshop on Causal Inference, Graphical Models and Prediction in honour of A. Philip Dawid, University of Cambridge, 30 April.

MAY

LAMAS: 7th Workshop on Logical Aspects of Multi-Agent Systems, Paris, France, 5–6 May.

MSDM: Workshop on Multi-Agent Sequential Decision Making Under Uncertainty, Paris, France, 5–6 May.

SQUARE: 4th World Congress on the Square of Opposition, Pontifical Lateran University, Vatican, 5–9 May.

ADMI: 10th International Workshop on Agents and Data Mining Interaction, Paris, France, 5–9 May.

MS6: Models and Simulations 6, University of Notre Dame, 9–11 May.

EIDYN: Normativity and Modality, Edinburgh, 9–11 May.

FORMAL METHODS: Singapore, 12–16 May.

WPI: 6th Workshop in the Philosophy of Information, Duke University, 15–16 May.

SLACCR: St. Louis Annual Conference on Reasons and Rationality, St. Louis, MO, 18–20 May.

SCIENCE & METAPHYSICS: Ghent, Belgium, 20–21 May.

ABSTRACTION: Philosophy and Mathematics, Oslo, 21–23 May.

WFAP: Language and Philosophical Method, University of Vienna, 22–24 May.

ARGDIA: 12th ArgDiaP Conference “From Real Data to Argument Mining”, Warsaw, Poland, 23–24 May.

MAP: Mathematics, Algorithms and Proofs, Paris, France, 26–30 May.

FILMAT: 1st International Conference of the Italian Network for the Philosophy of Mathematics, Milan, 29–31 May.

FORMAL ETHICS: EIPE, Erasmus University Rotterdam, 30–31 May.

JUNE

MSLP: Mathematising Science, University of East Anglia, 1–3 June.

F& MI: Fundamentality and Metaphysical Informatism, University of Helsinki, Finland, 2–3 June.

ALGMATHLOG: Algebra and Mathematical Logic: Theory and Applications, Kazan, 2–6 June.

CWAP: Normativity of Meaning, Belief and Knowledge, Krakow, Poland, 4–6 June.

LOGICMATHPHYSICS: Ontario, Canada, 5–6 June.

TECHNOCOG: Innovation and Scientific Practice, Barcelona, 5–6 June.

POP: 4th LSE Graduate Conference in Philosophy of Probability, London, 6–7 June.

LG& M: Logic, Grammar, and Meaning, University of East Anglia, 7–9 June.

EC: 15th ACM Conference on Economics and Computation, Stanford University, CA, USA, 8–12 June.

MoT: Truthmaking as Grounding: For and Against, Barcelona, 9–10 June.

CCR: 9th International Conference on Computability, Complexity and Randomness, Singapore, 9–13 June.

PARACONSISTENCY: Paraconsistent Reasoning in Science and Mathematics, Munich, Germany, 11–13 June.

IYSM: International Young Statistician Meeting, Università di Cagliari, Italy, 13–14 June.

COLT: 27th Annual Conference on Learning Theory, Barcelona, 13–15 June.

LOGICA: Hejnice, Czech Republic, 16–20 June.

SILFS: International Conference of the Italian Society for Logic and Philosophy of Sciences, University of Rome “Roma TRE”, 18–20 June.

AMSTA: 8th International KES Conference on Agents and Multi-agent Systems—Technologies & Applications, Crete, Greece, 18–20 June.

FEW: 11th Annual Formal Epistemology Workshop, University of Southern California, Los Angeles, CA, 20–22 June.

SEP: 42nd Annual Meeting of the Society for Exact Philosophy, California Institute of Technology, Pasadena, CA, 22–24 June.

CE: Computability in Europe, Budapest, Hungary, 23–27 June.


SPS: Metaphysics of Science, Lille, 25–27 June.

A & N: The “Artificial” and the “Natural” in the Life Sciences, University of Exeter, 25–27 June.

CogSciJR: Jagiellonian-Rutgers Conference in Cognitive Science, Kraków, Poland, 25–29 June.

SPE: Semantics and Philosophy in Europe, Berlin, 26–28 June.

&HPS: Integrated History and Philosophy of Science, Vienna, Austria, 26–29 June.



xkcd.com

EGEC: 4th Annual Edinburgh Graduate Epistemology Conference, University of Edinburgh, 27–28 June.

IPSP: Imprecise Probabilities in Statistics and Philosophy, LMU Munich, 27–28 June.

EEN: European Epistemology Network Meeting, Madrid, 30 June–2 July.

JULY

IACAP: Annual Meeting of the International Association for Computing and Philosophy, Thessaloniki, Greece, 2–4 July.

WCT: workshop on Computability Theory, Prague, 3–4 July.

OPEN MINDS: University of Manchester, 4 July.

SotFoM: Symposium on the Foundations of Mathematics, Kurt Gödel Research Center, University of Vienna, 7–8 July.

CICM: Intelligent Computer Mathematics, University of Coimbra, Portugal, 7–11 July.

TLXIV: Trends in Logic, Ghent University, Belgium, 8–11 July.

FLoC: 6th Federated Logic Conference, Vienna, 9–24 July.

BSPS: British Society for the Philosophy of Science, University of Cambridge, 10–11 July.

SIS: Scientific Meeting of the Italian Statistical Society, Cagliari, Italy, 11–13 July.

DEON: 12th International Conference on Deontic Logic and Normative Systems, Ghent, Belgium, 12–15 July.

SAT: 17th International Conference on Theory and Applications of Satisfiability Testing, Vienna, Austria, 14–17 July.

IPMU: 15th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Montpellier, France, 15–19 July.

LATD: Logic, Algebra, and Truth Degrees, Vienna, 16–19 July.

NMR: 15th International Workshop on Non-Monotonic Reasoning, Vienna, Austria, 17–19 July.

IJCAR: 7th International Joint Conference on Automated Reasoning, Vienna, Austria, 19–22 July.

AUAI: Uncertainty in Artificial Intelligence Conference, Quebec, Canada, 23–27 July.

KRC: Reasoning Conference, Konstanz, Germany, 24–27 July.

LOFT: Eleventh Conference on Logic and the Foundations of Game and Decision Theory, University of Bergen, Norway, 27–30 July.

UCM: Uncertainty in Computer Models 2014, University of Sheffield, 28–30 July.

AUGUST

AIML: Advances in Modal Logic, University of Groningen, 5–8 August.

ECAI: 21st European Conference on Artificial Intelligence, Prague, Czech Republic, 18–22 August.

ROBO-PHILOSOPHY: Aarhus University, Denmark, 20–23 August.

HYP0: Hypothetical Reasoning, Tübingen, Germany, 23–24 August.

SLS: 9th Scandinavian Logic Symposium, University of Tampere, Finland, 25–27 August.

ECAP: 8th European Conference of Analytic Philosophy, University of Bucharest, Romania, 28 August–2 September.

SEPTEMBER

WPMIIP: 7th Workshop on Principles and Methods of Statistical Inference with Interval Probability, Ghent, Belgium, 8–12 September.

COMMA: 5th International Conference on Computational Models of Argument, Scottish Highlands, 9–12 September.

ENPOSS: 3rd European Network for the Philosophy of the Social Sciences Conference, Madrid, 10–12 September.

GANDALF: 5th International Symposium on Games, Automata, Logics and Formal Verification, Verona, Italy, 10–12 September.

LANCOG: workshop on Modal Syllogistic, Lisbon, 11–13 September.

SCLC: 10th Symposium for Cognition, Logic and Communication, University of Latvia, Riga, 12–13 September.

NoR&N: Nature of Rules and Normativity, Prague, Czech Republic, 17–19 September.

IWSBP: 11th International Workshop on Boolean Problems, Freiberg, Germany, 17–19 September.

ICTCS: 15th Italian Conference on Theoretical Computer Science, Perugia, Italy, 17–19 September.

FOIS: 8th International Conference on Formal Ontology in Information Systems, Rio de Janeiro, 22–25 September.

KI: 37th German Conference on Artificial Intelligence, Stuttgart, 22–26 September.

JELIA: 14th European Conference on Logics in Artificial Intelligence, Madeira Island, Portugal, 24–26 September.

LANCOG: workshop on Analyticity, Lisbon, 25–26 September.

BELIEF: 3rd International Conference on Belief Functions, Oxford, 26–28 September.

COURSES AND PROGRAMMES

Courses

MLSS: Machine Learning Summer School, Reykjavik, Iceland, 25 April–4 May.

EPISTEMIC GAME THEORY: EPICENTER, Maastricht University, 12–23 May.

IGSAR: 2nd Interdisciplinary Graduate School on Argumentation and Rhetoric “Corpus Analysis in Argument Studies”, Polish National Academy of Sciences, Warsaw, 21–24 May.

NASSLLI: 6th North American Summer School in Logic, Language and Information, University of Maryland, College Park, 21–29 June.

EASLLC: 3rd East-Asian School on Logic, Language and Computation, Tsinghua University, China, 2–8 July.

CARNEGIE MELLON: Summer School in Logic and Formal Epistemology, 2–20 July.

SIPTA: 6th SIPTA School on Imprecise Probabilities, Montpellier, France, 21–25 July.

MCMP: MCMP Summer School on Mathematical Philosophy for Female Students, Munich, Germany, 27 July–2 August.

ESSLLI: 26th European Summer School in Logic, Language and Information, University of Tübingen, Germany, 18–22 August.

CLPA: Summer School on Argumentation: Computational and Linguistic Perspectives on Argumentation, University of Dundee, Scotland, 4–8 September.

Programmes

APhil: MA/PhD in Analytic Philosophy, University of Barcelona.

MASTER PROGRAMME: MA in Pure and Applied Logic, University of Barcelona.

DOCTORAL PROGRAMME IN PHILOSOPHY: Language, Mind and Practice, Department of Philosophy, University of Zurich, Switzerland.

HPSM: MA in the History and Philosophy of Science and Medicine, Durham University.

MASTER PROGRAMME: in Statistics, University College Dublin.

LoPhiSC: Master in Logic, Philosophy of Science & Epistemology, Pantheon-Sorbonne University (Paris 1) and Paris-Sorbonne University (Paris 4).

MASTER PROGRAMME: in Artificial Intelligence, Radboud University Nijmegen, the Netherlands.

MASTER PROGRAMME: Philosophy and Economics, Institute of Philosophy, University of Bayreuth.

MA IN COGNITIVE SCIENCE: School of Politics, International Studies and Philosophy, Queen's University Belfast.

MA IN LOGIC AND THE PHILOSOPHY OF MATHEMATICS: Department of Philosophy, University of Bristol.

MA PROGRAMMES: in Philosophy of Science, University of Leeds.

MA IN LOGIC AND PHILOSOPHY OF SCIENCE: Faculty of Philosophy, Philosophy of Science and Study of Religion, LMU Munich.

MA IN LOGIC AND THEORY OF SCIENCE: Department of Logic of the Eotvos Lorand University, Budapest, Hungary.

MA IN METAPHYSICS, LANGUAGE, AND MIND: Department of Philosophy, University of Liverpool.

MA IN MIND, BRAIN AND LEARNING: Westminster Institute of Education, Oxford Brookes University.

MA IN PHILOSOPHY: by research, Tilburg University.

MA IN PHILOSOPHY, SCIENCE AND SOCIETY: TiLPS, Tilburg University.

MA IN PHILOSOPHY OF BIOLOGICAL AND COGNITIVE SCIENCES: Department of Philosophy, University of Bristol.

MA IN RHETORIC: School of Journalism, Media and Communication, University of Central Lancashire.

MA PROGRAMMES: in Philosophy of Language and Linguistics, and Philosophy of Mind and Psychology, University of Birmingham.

MRES IN COGNITIVE SCIENCE AND HUMANITIES: LANGUAGE, COMMUNICATION AND ORGANIZATION: Institute for Logic, Cognition, Language, and Information, University of the Basque Country, Donostia, San Sebastián.

MRES IN METHODS AND PRACTICES OF PHILOSOPHICAL RESEARCH: Northern Institute of Philosophy, University of Aberdeen.

MSC IN APPLIED STATISTICS: Department of Economics, Mathematics and Statistics, Birkbeck, University of London.

MSC IN APPLIED STATISTICS AND DATAMINING: School of Mathematics and Statistics, University of St Andrews.

MSC IN ARTIFICIAL INTELLIGENCE: Faculty of Engineering, University of Leeds.

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

MSC IN COGNITIVE SCIENCE: University of Osnabrück, Germany.

MSC IN COGNITIVE PSYCHOLOGY/NEUROPSYCHOLOGY: School of Psychology, University of Kent.

MSC IN LOGIC: Institute for Logic, Language and Computation, University of Amsterdam.

MSC IN MIND, LANGUAGE & EMBODIED COGNITION: School of Philosophy, Psychology and Language Sciences, University of Edinburgh.

MSC IN PHILOSOPHY OF SCIENCE, TECHNOLOGY AND SOCIETY: University of Twente, The Netherlands.

MRES IN COGNITIVE SCIENCE AND HUMANITIES: LANGUAGE, COMMUNICATION AND ORGANIZATION: Institute for Logic, Cognition, Language, and Information, University of the Basque Country (Donostia San Sebastián).

OPEN MIND: International School of Advanced Studies in Cognitive Sciences, University of Bucharest.

PhD SCHOOL: in Statistics, Padua University.

JOBS AND STUDENTSHIPS

Jobs

POST-DOC POSITION: in Set Theory, Torino University, until filled.

POST-DOC POSITION: in Proof Theory, Faculty of Mathematics, Vienna University of Technology, deadline 7 March.

POST-DOC POSITION: on the project “Effective Bayesian Modelling with Knowledge Before Data”, School of Electronic Engineering and Computer Science, Queen Mary, University of London, deadline 21 March.

PROFESSOR: in Analytic Philosophy, Vienna, deadline 31 March.

POST-DOC POSITION: on the project “The Epistemology of Data-Intensive Science”, Egenis, University of Exeter, deadline 29 April.

Studentships

PhD POSITION: in Practical Philosophy, University of Lund, deadline 3 March.

PhD POSITION: on the project “Computational Model of Negotiation Skills in Virtual Artificial Agents”, School of Computing and Mathematics, University of Plymouth, deadline 10 March.

PhD POSITION: on the project “The Epistemology of Data-Intensive Science”, Egenis, University of Exeter, deadline 29 April.

PhD POSITION: on the project “Influence in Cyberspace: The relationship between information provenance, trust and identity within the context of cyber influence”, Web Science, University of Southampton, deadline 30 September.

MA IN REASONING

A programme at the University of Kent, Canterbury, UK. Gain the philosophical background required for a PhD in this area.

Optional modules available from Psychology, Computing, Statistics, Social Policy, Law, Biosciences and History.