
THE REASONER

VOLUME 16, NUMBER 8
NOVEMBER 2022

thereasoner.org
ISSN 1757-0522

CONTENTS


Editorial	66
Features	66
Dissemination Corner	66
News	69
What's Hot in ...	70
Events	71

EDITORIAL

Welcome to the last issue of The Reasoner for 2022.

It is a good occasion to remind you how you can get involved in it. This year we inaugurated a new format which we term *Focussed issue* – a collection of solicited contributions highlighting the state of the art of a field. If you'd like to coordinate one, please let us know. Topics covered so far include Evidential Pluralism, History of Logical Reasoning and Infinitary Reasoning.

Our *Dissemination corner* is a relatively recent feature which has been popular in the last couple of years. It allows the wider reasoning community to keep up with the developments of reasoning-related projects, read about your latest books, pa-



pers, open positions, and hopefully get inspired for the next successful reasoning-related project.

Last, but certainly not least, you may contribute by interviewing someone in your own field or by contributing regular “What’s hot in...” columns.

See <https://blogs.kent.ac.uk/thereasoner/submit/> for detailed information on how to submit.

HYKEL HOSNI

University of Milan

FEATURES

DISSEMINATION CORNER

Epistemic Risk and the Demands of Rationality

An election is about to take place, and you and I disagree about how likely it is that the ruling party will be returned to power. You're 90% sure they will, while I'm only 65% sure. Given our disagreement, must at least one of us be irrational? Surely not. After all, we might have different evidence—you might have access to polling data I haven't seen. But what if we both have the same evidence; in that case, can we disagree while both being rational? This is a particular instance of the question that divides permissivists—who think we can—from adherents of the Uniqueness Thesis—who think we can't. In *Epistemic Risk and the Demands of Rationality*, I argue in favour of permissivism. Indeed, I defend a radical version of permissivism about rational credence, which says that there is a wide range of rationally permissible prior credences, and therefore, for nearly every body of evidence, a



wide range of posterior credences that are rationally permissible for individuals who possess that evidence.

The conclusion has implications in many areas: it may scupper any hopes of a wholly objective science; it might affect how we respond when, in a criminal trial, we are asked whether what the defendant claims to have believed when they acted is something that a reasonable person would have believed in their situation; and, as I discuss in the final chapter of the book, it might affect how we understand conspiracy theorists and the wrongs they do to others when they act on their conspiracist beliefs.

My argument for this conclusion takes place within a teleological conception of epistemic rationality. That is, I assume that facts about what is epistemically rational depend on facts about what is epistemically good. I work within accuracy-first epistemology, which takes a veritist conception of what is epistemically good: a credence is better the more accurate it is; so, for instance, a credence in a true proposition is better the higher it is, while a credence in a falsehood is better the lower it is. And indeed, I follow much of accuracy-first epistemology in endorsing a particular class of measures of the accuracy of a set of credences – they are the strictly proper scoring rules (L. J. Savage 1971: “Elicitation of Personal Probabilities and Expectations.” *Journal of the American Statistical Association* 66.336, 783–801).

How do we determine what is epistemically rational on the basis of this conception of what is epistemically good? For this, we turn to decision theory. This might seem surprising, but decision theory is really a collection of bridge principles that allow us to determine which means it is rational to use to pursue our ends. In its standard application, that end is pragmatic utility, and it tells us which actions are rational means to take in our pursuit of that; in our case, the end is accuracy, and decision theory tells us which sets of credences are rational means towards that.

Take the norm of Dominance, for instance, which says that it is irrational to pick an option if there is another that is guaranteed to be better. In the context of accuracy-first epistemology, it says it's irrational to have one set of credences when there is another that is guaranteed to be more accurate. And, due to a result that has become something of a cornerstone of the approach, we can use this to justify the norm of Probabilism. For it turns out that, if you measure accuracy using a strictly proper scoring rule, and if your credences do not obey the probability axioms, then there are alternative credences, defined over the same propositions, that are guaranteed to be more accurate (J. Predd et al. “Probabilistic Coherence and Proper Scoring Rules.” *IEEE Transactions of Information Theory* 55.10 4786–4792.)

So the next stage of the book's argument seeks further decision principles, like Dominance, that might govern our choice of prior credences. It is at this point that the notion of epistemic risk enters. It is inspired by William James's position in William (1897: “The Will to Believe”. *The Will to Believe, and other essays in popular philosophy*. New York: Longmans Green). James was concerned with categorical beliefs rather than credences, but his position generalises to the credal case. He notes that, to adopt a belief is a risky business. You risk believing falsely, which is bad, but you also open up the possibility of believing truly, which is good. If, instead, you suspend judgment, then you are guaranteed a neutral outcome, neither good nor bad. James thinks that our attitudes to epistemic risk—whether we are risk-inclined and so in need of less evi-

dence before taking the risk of forming a belief, or risk-averse and so in need of more evidence before doing so—are subjective aspects of ourselves. He is permissivist about epistemic risk: just as it is rationally permissible for you to love Brahms' music and for me to loathe it, so it is rational for you to be epistemically risk-inclined and for me to be epistemically risk-averse.

My book is an attempt to formalise James' view and build a coherent version of it. One of the themes that runs through it is how difficult it proves to accomplish that: we find ourselves having to steer between the Scylla of inconsistency and the Charybdis of impermissivism, often coming very close to peril.

The family of decision rules I formulate and for which I argue on the basis of James' permissive account of epistemic risk is called the Generalized Hurwicz Criterion. Each member of it corresponds to a different attitude to epistemic risk; some are more risk-inclined, others more risk-seeking. Having formulated them, I ask what each demands of an individual who uses it to pick their prior credences. The more risk-averse members of the family require you to pick the uniform prior, while the more risk-inclined members permit you to assign greater credence to some possibilities than to others. And from this our permissivism about rational credences follows.

RICHARD PETTIGREW
University of Bristol

Philosophy of Biomimetics

Biomimetics aims to use our knowledge of living nature to develop new technical products. But how exactly does that work? And can biomimetics be considered an independent scientific discipline? A new research project “Learning from Nature” (<https://biomimetics.hypotheses.org>) at the universities of Rostock and Tübingen is investigating these questions.

From bird flight to gecko feet, humans have always been inspired by nature for technical developments. In the 20th century, this has been elevated to a programme under the name of biomimetics, and it is hard to imagine our everyday life without biomimetic products such as the Velcro fastener. From the perspective of the philosophy of science, however, biomimetics has not yet been explored. Dr. Manfred Drack from the Institute for Evolution and Ecology at the University of Tübingen and Dr. Ludger Jansen from the Institute for Philosophy at the University of Rostock address this research gap and investigate the question of how exactly the transfer of biological knowledge to technical developments works. The research is funded by the German Research Foundation (DFG) with around 500,000 euros. The project will examine selected biomimetic development projects to see whether common methods and a uniform subject matter can be found in them, which would justify conceiving biomimetics as a unified scientific discipline. Doing so, the project will in particular focus on reasoning and justification processes going on in interdisciplinary biomimetic research projects. In addition, the basic categories identified in the research processes will be formalised on the basis of a philosophical analysis so that they can be processed by a computer.

The overall goal of the project is to deepen our understanding of the new developments at the crossroads of biology and technology by a thorough analysis and consolidation of the the-

oretical foundations of biomimetics. The survey of the state of the art revealed the following open questions:

- Can biomimetics be conceived as a coherent field of research? What is the nature of biomimetic knowledge, and how is it derived?
- What is the specific object of biomimetic research and how is it to be analysed ontologically?

Answering these questions will shed new light not only on the exciting developments at the intersections of biology and technology, but also contribute to the ontological analysis of the underlying reality in domains that have long been considered as separate realms. Our leading hypothesis is that biomimetics is unified by epistemological features including the goal of studying the transfer of working principles for functions from biological to technical constructions. We will research how knowledge about functions, working principles and constructions can be gained and analyse ontologically what kind of entities they are and how they fit into the wider domain of biomimetics. This will contribute new insights to the debate about functions, adding an interdisciplinary perspective and an intertwined analysis of working principles and constructions. We expect that our results will also foster practical communication in biomimetics as well as help lay the basis to systematise existing knowledge and to generate new knowledge in the field.

In order to accomplish the above-mentioned objectives, we plan to

- integrate the project within the wider biomimetics community and compose a corpus of research examples to be analysed in the project,
- analyse whether it is justified to view biomimetics as a unified field of research with its specific set of epistemological features,
- analyse reasoning patterns in biomimetic projects,
- review the ontological commitments of biomimetics to evaluate existing approaches,
- analyse the foundational categories of biomimetics,
- evaluate these results through presentation and discussion at an expert workshop,
- formalise the ontology,
- and evaluate it for consistency and adequacy.

To enhance communication on the philosophy of biomimetics, the project has already set up a website (<https://biomimetics.hypotheses.org>) and a mailing list (<https://listserv.uni-tuebingen.de/mailman/listinfo/biomimetics-philosophy>), to which all interested researchers and practitioners are invited. In the near future, a regular series of online-events is planned with presentations and workshop discussions on theoretical aspects of biomimetics.

MANFRED DRACK

Eberhards Karls University Tübingen

LUDGER JANSEN

University of Rostock

BRIO

B.R.I.O. (Bias, Risk and Opacity) is a national research project (PRIN) funded by the Italian Ministry of University and Research (<https://sites.unimi.it/brio/>). The project focuses on developing design criteria for Trustworthy AI based on philosophical analyses of transparency, bias and risk combined with their formalisation and technical implementation.



The first research meeting of the project took place on September 17th in Milan and represented an occasion for internal scholars to share their initial results, as well as for external researchers to present their works and look for potential collaborations. The research meeting lasted the whole day and was divided in one morning and one afternoon session. The morning session hosted an opening lecture by the guest speaker prof. Stephanie Dick from Simon Fraser University School of Communication. For the rest, it was dedicated to general presentations by investigators of the different research units. The afternoon session, on the other hand, was structured as a poster session with contributions from young researchers both internal and external to the project.

In the opening lecture, prof. Stephanie Dick proposed a reconstruction of the history of AI in the United States from the end of the second world war to nowadays. The proposed reconstruction identifies a progressive historical shift from an AI conceived as a mere mathematical tool for automatic theorem proving to a “social” and “embedded” AI that assists humans in everyday life. A direct consequence of this shift is the emergence of a dynamic of trust between humans and AI that raises new fundamental issues and risks related to opacity and bias. Addressing these issues is going to represent the major challenge for AI research in the near future.

The first talk “Proof-checking bias for labelling methods”, given by Fabio D’Asaro (UniVR), concerned a typed natural deduction system developed and designed to formally verify the presence of “bias” in automatic labelling methods. The talk was based on a recent paper by D’Asaro and Primiero (2022: Checking trustworthiness of probabilistic computations in a typed natural deduction system, <https://arxiv.org/abs/2206.12934>, CoRR abs/2206.12934) and offered an innovative interpretation of bias as a measure of the divergence between the expected probabilistic labelling by a classifier trained on opaque data and the fairness constraints set by a transparent reference data-set.


The second talk “Formula Understandability in Description Logic”, given by Daniele Porello (UniGE), concerned applications of description logic to eXplainable AI. The talk focused specifically on the *Tooth operator*, a weighted threshold operator recently introduced in description logic. The expressions formulated using this operator (Tooth expressions) incorporate some insights stemming from the cognitive-oriented theories of concepts, such as the prototypes theory. A recent study by Righetti et al. (2022: Evaluating the interpretability of threshold operators. In *International Conference on Knowledge Engineering and Knowledge Management*, pages 136-151) has proved that Tooth expressions are easier to understand for users

than equivalent expressions formulated in the standard disjunctive normal form. This, Porello claimed, strongly suggests the potential usefulness of description logic with Tooth operator to construct local explanations of black box classifiers in the context of Explainable AI.

The third talk “The Reproducibility Crisis and its Impact on TAI”, given by Viola Schiaffonati (PoliMI), focused on the impact of reproducibility crisis on trustworthy AI. Notably, it outlined some basic ideas and guidelines for what is meant to be a future broad-ranging epistemological investigation of the reproducibility crisis in AI. The speaker stressed in particular the fact that reproducibility is not just a matter of openness and cannot be solved by simply implementing openness policies. On the contrary, limitations in reproducibility are an intrinsic and usually unavoidable characteristic of many contemporary AI technologies. What we want is such limitations not to eventually translate into a crisis of experts’ confidence in AI methods and technologies. For this to happen, Schiaffonati claimed, a cultural shift is necessary that put into question the status of reproducibility as a gold-standard and opens to other ways for ensuring the reliability of AI systems.

Posters in the afternoon session included various contributions from different technical and philosophical research areas related to the project. Contributions in the area of epistemology and ethics of AI focused mostly on data quality dimensions (Canali and Quaresimi) and the ethics of recommender systems (Pedrazzoli). Contributions in the area of applied logic, on the other hand, concerned relational semantics (Kubyshkina) and model checking methods (Termine) for evaluating trustworthiness of AI systems. Finally, contributions in the area of eXplainable AI included: bias identification (Manganini), robustness metrics for deep learning (Posipillo), methods to measure risk awareness (Alnazer) and mathematical tools to build explainable neural networks (Bocchi). All the posters are available on the web page of the BRIO project dedicated to the event (URL: <https://sites.unimi.it/brio/first-brio-research-meeting>).

ALBERTO TERMINE
LUCI Group, University of Milan



MY NEW VACATION SPOT HAS VERY FAST INTERNET THAT TURNS OFF RANDOMLY EVERY NOW AND THEN, JUST SO YOU CAN TELL PEOPLE YOU'LL BE STAYING SOMEWHERE WITHOUT A RELIABLE CONNECTION.

Philosophy & Finance Network (‘Phinance’) Biennial Conference, 8-9 September 2022

Phinance’s first biennial conference provided an opportunity for the growing philosophy of the financial community to take stock of several important issues in philosophy and finance.

Ekaterina Svetlova (Twente) reviewed AI’s ethical responsibility in finance. Moving from the fact that AI actors in financial markets are often *ignorant* about how AI systems work, she examined under what conditions these AI actors can be held responsible for their ignorance.

Paolo Barucca (UCL) dealt with AI and the efficient market hypothesis (EMH). Noting that markets and AI algorithms are complex computing systems which take information (fundamentals) and return more information (prices), he asked if AI algorithms improve markets efficiency. He noted that it depends on how these algorithms are adjusted but that they are not able to choose economic priorities and rules for us.

Christian Walter (Paris) focused on two mathematical representations of the EMH: the mean-variance under the *real-world* probability P and the martingale pricing under the *risk-neutral* probability Q. Noting that these two representations can be characterized as two quantification conventions, he argued that they play a *performative* role.

Boudewijn de Bruin (Groningen) approached the urgency of taking drastic measures to avert climate change and looked at what sustainable finance could do about it. He looked at the *EU 2018 Action Plan: Financing* and links it to philosophy and ethics by asking seven conceptual questions about the roles and responsibilities of finance—e.g. the link between financial regulation and climate change mitigation.

Francesco Guala (Milano Statale) presented new insights into how people conceptualize money. As the money is considered as such in virtue of (1) physical properties, intended function and history, (2) the collective recognition, and (3) current use and function, he provided evidence that a many people take ‘recognition’ as sufficient for money to exist, but not ‘history’ alone, and consider ‘function’ to be important.

Taylor Spears (Edinburgh) examined the translation of a set of models and associated practices from one field to another. He made two points. (1) The development trajectories of models are shaped when they function as market infrastructures. (2) Two forms of infrastructural alignment, i.e. *ontological* and *material*, must be met for models to become taken-for-granted infrastructures.

Roberto Violi (Bank of Italy) examined the implication for regulatory capital requirement resulting from the risk model and risk classes aggregation in the Standardized Approach (SA) and Simplified Standard Approach (SSA). Since some significant issues can produce overestimates of capital requirements in SSA, he argued that this cannot be justified through a generic search for a prudent regulatory framework for the bank’s capital requirement.

Melissa Vergara-Fernandez (Erasmus - Rotterdam) focused on the empirical turn in economics. She compared Backhouse & Cherrier’s general story of economic change with that of the financial economy (asset pricing). She argued that based on their similarities, the latter field may be empirically successful because of its empirical turn.

Emiliano Ippoliti (Sapienza) examined how ‘micro-quotes’

undermine the epistemic value of quotes and the core functions of financial systems. He argued that the choice of the target and time threshold of speed bumps (a solution to micro-quotes), is not a ‘merely technical’ choice: rather it has significant impact on the preferred structure & agents of financial markets and is epistemically and morally laden.

EMILIANO IPPOLITI
Sapienza University of Rome

Logic for the new AI spring, 12-16 September 2022

The *Logic for the new AI Spring* school was held at the lake Como (Italy) and aimed at bringing together scholars working within the AI, and underlining the fundamental role of logic in AI research and its applications. The school consisted of four tutorials, work in progress sessions, and a general lecture by Michael Wooldridge (Oxford University).



The school started with the tutorial on multiagents systems given by Michael Wooldridge. The tutorial presented a survey on cooperation logics in game-like multiagent systems; the normative systems representing social laws with the final goal of searching effective strategies for achieving desirable objectives; Boolean games, with a particular interest on calculating epistemic costs for agent’s choices and their regulations; rational verification, that is, the problem of checking whether the system exhibits some specific behaviour under the assumption that agents within the system act rationally. The material of the tutorials provided Wooldridge with a solid ground for his final general lecture on understanding the functioning of AI and discussing its limits.

Stephanie Dick (Simon Fraser University) gave a tutorial on the history and culture of. The tutorial provided material on the early history of AI by contextualising the development of automated calculations within specific social and political conditions; a discussion on what it means for AI to be “human” and “to reason”; the involvement of knowledge and interaction in human-machine collaboration; application, applicability, and limits of AI in solving scientific and social problems; open questions concerning possible definitions of “fairness”, “accountability”, and “transparency” for AI.

The third tutorial given by Josef Urban (Czech Technical University in Prague) was about machine learning and theorem proving. The tutorial presented the historical discussion on learning and reasoning in mathematics and in science; combination of learning and reasoning; the foundations of theorem provers (set theory, higher-order logic, type theory, and other logical frameworks) and their implementations; the main learning approaches (the data- and theory-driven methodologies) with various examples and discussion on features and data preprocessing; the questions of high-, low-, and mid-level reasoning guidance for theorem proving; the problem (and existing solutions) of synthesizing math expressions in AI systems and autoformalization of AI.

The last tutorial, entitled “Logical Foundations of Categorization Theory”, was given by Alessandra Palmigiano (Vrije

Universiteit Amsterdam). The tutorial discussed the interdisciplinary character of studies in categorization; Lattice Expansions logics, their intuitive meaning, and the semantic frameworks of these logics; various views on categorization (classical, prototype, exemplar, and theory-based views), their advantages and disadvantages; conceptual spaces, which permit one to reason about objects and their categorization, and a logical representation of typicality, which is one of the central ingredients for categorizing objects.

An important part of the school was devoted to the presentation of work in progress related to the studies in AI by Master students, Ph.D. candidates and postdoctoral researches from various institutions in Italy and abroad. The range of topics was quite broad from machine learning to neurosciences, linguistics, and epistemic logic.

The school belongs to H2020 the project MOSAIC (<https://sites.google.com/view/mosaic-rise/>) with funding from the Fondazione Volta, the Department of Philosophy of the University of Milan, the project BRIO (<https://sites.unimi.it/brio/>), the Turing Center at ETH Zurich, and the History and Philosophy of Computing Commission (HaPoC),

A special issue on the *International Journal of Approximate Reasoning* is attached to the event. For more details on the submission see the journal’s website (<https://www.sciencedirect.com/journal/international-journal-of-approximate-reasoning/about/forthcoming-special-issues>).

EKATERINA KUBYSHKINA
LUCI Group, University of Milan

Calls for Papers

TEMPORAL REASONING AND TENSED TRUTHS: special issues of *Synthese*, deadline 31 December.

THE METAPHYSICS AND SEMANTICS OF TRUTHMAKING: special issue of *The Philosophy of Science*, deadline 31 December.

LOGIC FOR THE NEW AI SPRING: special issue of *International Journal of Approximate Reasoning*, deadline 1 March.

CAUSAL INFERENCE, PROBABILITY THEORY AND GRAPHICAL CONCEPTS: special issue of *Computation*, deadline 30 April.

WHAT’S HOT IN . . .

Mathematical Philosophy

Philosophers in the 2000s have been stirring from their armchairs and turning increasingly toward *practice-based* topics and methods. In philosophy of empirical science, this has looked like an abating obsession with laws, theories and the realism question, and a corresponding swell of interest in models, experiments, data, measurement, imagination, and so on. In philosophy of mathematics, we’ve begun to acknowledge the existence of subjects beyond



set theory, arithmetic and geometry, and of epistemic goals beyond proof.

At least two insights lie behind these shifts in fashion. First, taking practice seriously helps us ask new and better *questions*: we couldn't have had an enlightening debate about explanation in mathematics, for example, if we'd never deigned to notice how seriously mathematicians themselves take explanatory concerns. Second, the practical turn is useful because it gives us access to new and better *evidence*: if I want to know whether the Bohr model has advanced science's understanding of the atom, it's much more informative to consult the judgments of working physicists than to query my own half-baked intuitions.

So the practical turn has, I think, done lots of good. But you might wonder whether its scope should be wider still. (I've sometimes thought, for instance, that aesthetics could benefit from greater engagement with artistic practice. But let's confine ourselves to *Reasoner* territory for now.) Are we sure we've applied these methods everywhere they can be fruitfully applied?

Ben Martin thinks not. In "[The philosophy of logical practice](#)" (*Metaphilosophy*, 2022), Martin makes a case for logic as a domain deserving its own practice-based tradition, as "a new field of research... to sit alongside traditional philosophy of logic" (268). (As he notes, earlier anticipations of this program appear in Dutilh Novaes's "[Towards a practice-based philosophy of logic: Formal languages as a case study](#)" (*Philosophia Scientiae*, 2012) and Payette and Wyatt's "[How do logics explain?](#)" (*Australasian Journal of Philosophy*, 2018).)

Martin's motivation for PLP has much in common with the standard justifications for philosophy of scientific and mathematical practice. He lists five grievances with traditional philosophy of logic: (1) Its picture of logic is too simplistic and idealized; (2) It focuses too much on metaphysical issues imported from mainstream philosophy, e.g. about the ontology of logic and the nature of logical facts; (3) It privileges logics and applications that philosophers happen to find interesting; (4) It ignores the history of logic and wrongly treats current presumptions as inevitable; (5) It treats logic as a set of finished formal systems, disregarding the goals, methods and standards that led to their development.

This is a good list; anyone with a mite of sympathy for the approach will find it easy to nod along. But it would also be nice to see what PLP looks like in action. What sorts of issues will it take on? What conclusions might it yield? How if at all does it differ from philosophy of scientific and mathematical practice—and can other practice-oriented philosophers learn anything from those differences?

Fortunately, Martin has the reader's back. Section 4.2 of his paper delivers a sketch of a case study. Its subject is our knowledge of validity: how do we come to accept a given (rule of) inference in a candidate logic as deductively correct?

Mainstream philosophers of logic have typically shown up to this party with commitments to defend and axes to grind. For instance, they generally assume that we know some facts about validity, and that these facts must be known a priori and non-inferentially. The question is then how to make sense of this epistemology within one's preferred philosophical framework. Rationalists are happy to posit a faculty by which we directly apprehend valid inferences. This won't fly with empiricists, who are thus compelled to identify knowledge of validity with linguistic competence. In Martin's telling, this "top-down" approach has led to a long and fruitless battle between two ulti-

mately unappealing views.

The problem with the mainstream debate, from the practice-based philosopher's viewpoint, is that neither side has bothered to consider the process by which logicians actually formulate, test and defend prospective inference rules. This oversight should worry us. If anyone has reliable methods for gaining logical knowledge, it's surely the experts who have thought longest and hardest. And if their methods don't involve appeals to a priori rational insight or linguistic conventions, we should question the wisdom of the traditional debate.

In fact, as Martin points out, logicians appeal to many forms of evidence when weighing prospective inference rules, such as their mathematical fruitfulness and their utility for solving puzzles and defusing paradoxes. Neither rationalists nor semanticists have a clear story to tell about why these factors should be epistemically relevant. So either the philosophers have made mistaken assumptions, or the logicians are deeply confused about how to do their own work. PLP knows which side it's on.

It'll come as no surprise that I, a card-carrying philosopher of mathematical practice, am quite on board with all this (and the many more true and interesting things in Martin's paper, which you should go read). But I have a question or two.

Here's one. A strong argument in favor of practice-based philosophy of math, it seems to me, is that exceedingly few philosophers have research-caliber mathematical expertise, and only a small minority have any training beyond the undergraduate level. So there's just no reason to expect philosophers' intuitions to be any good, at least when it comes to questions of much importance or sophistication. By contrast, lots of philosophers have done serious logic, and more have a respectable background. So a determined traditionalist might deny that philosophers of logic are obliged to defer to full-time logicians, at least on many issues. (Indeed, as competent speakers and reasoners, one might think we all have as much insight as any expert into whether *P* implies *Q*.)

For the record, I don't endorse this argument myself! But it's one I'd like to see answered, if only for the greater glory of practice-based philosophy.

WILLIAM D'ALESSANDRO
MCMP Munich

EVENTS

NOVEMBER

PoP: Philosophy of Physics, virtual, 2 November.

EW&ER: Epistemic Wrongs and Epistemic Reparations, University of Johannesburg, 3–4 November.

C&E: Causation, Explanation, and Everything Else: Celebrating the Career of James F. Woodward, Pittsburgh, Pennsylvania, 4 November.

M-VL: Finite Model Theory and Many-Valued Logic, virtual, 10–12 November.

EIRMC: Epistemic Injustice in the Medical Context, University of California, Irvine, 17–18 November.

MARCH

HPS: Integrated History and Philosophy of Science, University of South Carolina, 16–18 March.

COURSES AND PROGRAMMES

Programmes

MA IN REASONING, ANALYSIS AND MODELLING: University of Milan, Italy.

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.

DOCTORAL PROGRAMME IN PHILOSOPHY: Department of Philosophy, University of Milan, Italy.

LOGICS: Joint doctoral program on Logical Methods in Computer Science, TU Wien, TU Graz, and JKU Linz, Austria.

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

LOPHISC: Master in Logic, Philosophy of Science and 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 OF BIOLOGICAL AND COGNITIVE SCIENCES: Department of Philosophy, University of Bristol.

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

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 SYSTEMS: Language, Learning, and Reasoning, University of Potsdam.

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.

RESEARCH MASTER IN PHILOSOPHY AND ECONOMICS: Erasmus University Rotterdam, The Netherlands.

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

MA IN PHILOSOPHY: Dept. of Philosophy, California State University Long Beach.

JOBS AND STUDENTSHIPS

Studentships

12 PHD GRANTS: in Mind, Brain and ReasoningThe Human Mind and Its Explanations: Language, Brain and Reasoning , University of Milan, deadline 2 September.


Jobs

JUNIOR PROFESSORSHIP: in Philosophy of Science, University of Dresden, open until filled.

POST-DOC: in Epistemic Logic, University of Bayreuth, Germany, deadline 18 November.

POST-DOC: in Philosophy of Science, University of Minnesota, deadline 2 December.

ASSISTANT PROFESSOR: in Philosophy of Science, London School of Economics, deadline 12 December.



The Human Brain:
Can do calculus, but is also reserving judgment about whether the guy in the mirror is a ghost.