
THE REASONER

VOLUME 11, NUMBER 6

JUNE 2017

thereasoner.org

ISSN 1757-0522

CONTENTS

Editorial

Features

News

The Reasoner Speculates

Reviews

What's Hot in ...

Events

Courses and Programmes

Jobs and Studentships

EIPE, those more interested in decision theory would bring up Peter Wakker.

26 Having enjoyed many years of reading THE REASONER, it is wonderful to have the honour of editing this issue. As a philosopher of economics based in Rotterdam, it was the perfect excuse to sit down with Peter Wakker and ask him about his intellectual biography and career, chess, experiments, the Archimedean axiom, and – of course, probability, statistics, and Savage. I hope you enjoy reading the interview as much as I did talking with Peter about these topics.

32 This issue also kick-starts the column ‘What’s hot in Economics & Philosophy?’, in which I will regularly review some trends in the burgeoning interdisciplinary field between philosophy and economics. This time around, it will focus on some things that have been going on at EIPE in Rotterdam – where the 20th anniversary celebrations of the institute have started in March 2017 with an international conference. The upcoming editions of the column will be a lot less inward-looking, though!

EDITORIAL

When I was about to move to Erasmus University Rotterdam several years ago, there were two topics of conversation that mentors and colleagues from all over the world would bring up. One of them was EIPE, the Erasmus Institute for Philosophy and Economics, which I was about to join, and its associated Research Master and PhD programme. The other topic was the work of decision theorist Peter Wakker and his research group. Those more interested in philosophy of economics and philosophy of science would bring up



CONRAD HEILMANN

Erasmus Institute for Philosophy and Economics (EIPE)
Erasmus University Rotterdam

FEATURES

Interview with Peter Wakker

Peter Wakker is a leading decision theorist, based at the Erasmus School of Economics at Erasmus University Rotterdam, working with a team of behavioural economists, which features, amongst others, Aurelien Baillon, Han Bleichrodt, and Kirsten Rohde: <https://www.eur.nl/ese/behec/>. Peter Wakker’s work comprises many crucial contributions to decision theory and statistics, both concerning theory and empirical work. He is perhaps most widely known for his ap-

plied work in behavioural economics, contributions to medical decision-making, and prospect theory (Wakker, P.P. (2010), “Prospect Theory: For Risk and Ambiguity”, CUP). His honours include the Medical Decision Making Career Achievement Award, the Frank P. Ramsey Medal from the INFORMS Decision Analysis Society and an Honorary Doctorate from the University of St. Gallen. His website harbours extensive ‘annotated references on decision and uncertainty’: <https://personal.eur.nl/wakker/>. I talked with Peter Wakker in March 2017 and would like to thank Hans Peters and Emanuele Di Francesco for their help in preparing this interview.

Conrad Heilmann: Peter, many thanks for taking the time to talk! Let’s start at the beginning. When did you first become interested in research?

Peter Wakker: As a six-year-old, I wanted to be a custom’s agent at the border, like my father, and then I wanted to be a veterinarian for a while. Around 12, I started to have this idea that I would become a researcher. I read books by famous scientists... – but I was not thinking about mathematics yet, that would be too abstract for a kid of 12. I was broadly interested in almost all disciplines of science, except economics (laughs)... – It was a close call: if it had not been mathematics, it would have been zoology. All my life I have been interested in that, too, reading books about it. But around age 15, I started to get into logical reasoning and that fascinated me. So then it became clear that mathematics fits me best, but I was in one of the last cohorts that had to take classes in everything. I took Greek and Latin, German, French, English, and all the natural sciences as well. I am really happy about that education. It was good, and there was no economics in there.



CH: You then studied mathematics as an undergraduate?

PW: I started studying a mix of mathematics and physics, but physics is horrible, you have to be in the laboratory... – you really have to work! In mathematics, you just sit in your chair, you have a cup of tea, you think a bit: and there it is. So, I quickly stopped physics, and it became only maths. (laughs) We had a course in quantum mechanics and I liked that and every other course that involved a probability concept. That’s how things came to life for me.

CH: What kind of statistics did you study?

PW: When I took courses in statistics, from the beginning, when I heard about classical statistics, I thought: this is just playing with numbers. I was born a Bayesian. At some stage, our teacher told us: ‘You cannot assign a probability to life on Mars, because either it exists or it doesn’t’, and that: ‘Frequency is the interpretation of probability’, and I immediately thought: ‘No, that doesn’t make sense!’ You have to balance what you know; you have to make a decision. At some stage I talked to the teacher and he said: ‘There is a crazy Italian who has such ideas’, and he wrote ‘de Finetti’ on

a piece of paper. With that piece of paper I went to the library, and I found his books. Reading them I thought: ‘This is what I will work on all my life.’

CH: And then you started to work on a PhD in decision theory?

PW: Well, it was more complicated than that. I really wanted to work on Bayesian decision theory, but nobody wanted to supervise me on this topic. Luckily, people in Leiden still took me in, because of my high grades and because I already had a paper published. So, they accepted me despite my stubborn attitude, but I was mostly on my own trying to find out what to work on. I even independently invented the convexity property for preferences. I have written, hand-written, hundreds of pages on all kinds of properties of convexity... You can find it in every economic textbook, but I didn’t have any idea about economics... – At some stage, however, I was getting desperate. Famous professors don’t want to discuss the foundations of the field with some unknown PhD student who is not well articulated. So that was a tough time and I thought that I’d probably fail. All my friends were going to conferences and getting published, and I didn’t know how to get that done. So, I thought I would probably fail and said to myself: ‘Well, then I’d become a teacher for secondary school’, and I got the degree for that.

CH: Yes, I saw on your CV that you have a secondary school teaching qualification – so this was really intended as a sort of professional lifeline?

PW: Yes, those were difficult, desperate times. But I was lucky: at a conference, I met Stef Tijs. [Stef Tijs, the ‘god-father of game theory in the Netherlands’, is amongst other contributions known for his ‘ τ -value’, see Tijs, S.H. (1987), [An axiomatization of the \$\tau\$ -value](#), *Math Soc Sci*, 13:177–81.] He had taught me in mathematics in Nijmegen and he told me: ‘Peter, those things that you are interested in, they are happening in economics, you should look into economics journals!’ Then I had to turn to economics.

CH: That’s quite a story! How, then, did your PhD come along?

PW: Stef Tijs became my supervisor, but he was not a full professor yet, so Pieter Ruys became my supervisor as well. I was doing a bit of game theory with Stef Tijs, and that was very nice. But for me, he was too much of a mathematician, and he didn’t see much relevance in other things, so I disagreed with him on that.

CH: Hans Peters told me about very lively discussions in the PhD seminars at the time between all of you...

PW: That’s right. But we never clashed; Stef Tijs is a wise person, and he let me do whatever I wanted. He introduced me to David Schmeidler at a conference. I visited him in Tel Aviv for six weeks as a PhD student, and then I worked on his model and did some theorems about that and he liked that. So that was part of my thesis.

CH: Not only in your PhD, but also afterwards, you have

worked on many different things – a lot of experimental work, all the contributions to prospect theory, and more foundational work.

PW: In my heart, it's theory. But I did lots of experimental work. Experiments are good: people who only work with mathematical models have never seen anything empirical, and they do not really understand what they are talking about. So if you also do the empirical work, you will understand what the concept means, and you'll connect to reality. My eight years of applied work in a hospital help a lot with that, too. [Wakker, P.P. (2008), *Lessons Learned by (from?) an Economist Working in Medical Decision Making*, *Med Decis Making* 28:690–8.] But, also, I was a little bit pragmatic.

CH: Pragmatic – in what way?

PW: Well, if you want to get the attention of people for your ideas, you must show to people that you are worth their attention; you must prove that. And people can only think that if you play their game. I can now work more closely on things that are really my interests, but I have all kinds of ideas about my biggest interest that I've not written yet and maybe I'll do it after some time. Right now, I need all my effort to keep my career, and the people and the group, to get some reputation. So, I need all my effort for that. If you have different ideas, it's not easy to sell them in a way in which the academic world and its rules allow for. So I had to publish papers in the way that journals accept papers. I was writing all kinds of papers that I didn't write with my whole heart.

CH: Looking back, what are the topics that are most 'close to your heart'? Do you also mean non-research things like chess?

PW: When I was around 18, I thought about becoming a professional chess player, but my competitors were always defeating me and I had to stop at age 25. Chess stays with me, though. I daydream of opening moves, and I keep reading and playing games of grandmasters, wanting to understand them. But now, I never play the game myself. Chess overlaps too much with my work. It's the same part of the brain. Playing chess is at the cost of my work. In my spare time, I prefer to do different things, like listening to music and cooking food.

CH: So, your one and only true calling remains Bayesianism?

PW: Indeed, my aim is basically to convert all humankind to become Bayesian: that is the goal of my life! And this includes all statisticians! (laughs) This was from the beginning my big interest and I couldn't believe that people could be doing other things, especially statisticians. But at this moment the behavioural approach is popular, and some people are unfortunately using it normatively. Ambiguity (unknown probability) now is really popular, and people are saying that policy-makers should be ambiguity averse. I do a lot of work on ambiguity, but I don't think it's rational to use any of these models normatively.

CH: There is a rumour that you once refused to co-author a paper if it would not identify Savage's decision theory as the

rational standard...

PW: I know what you are referring to. At some stage, I was about to co-author an introduction for a collection in honour of David Schmeidler. But I could never co-author a paper that explicitly says that deviations from Bayesianism can be rational; I will just never do that. And all my co-authors were very eager to write that. So, I told them from the beginning that it was not going to happen; it can happen, but with a footnote that I disagree. I know that as a co-author you need to be a bit flexible, but this has always been – and will always be – my position...

CH: Returning to your research plans, could you say a bit more about those things that you still want to do?

PW: I have in mind a sort of a theory, that would be more of a mathematical and philosophical theory, and that epitomises the normative state of the Bayesian model, and mostly the Sure Thing Principle in Savage's axioms. This has usually been the dividing line between all different kinds of theories and I want to go back to that.

CH: So, that would be a reformulation of Savage?

PW: Yes, but at a fundamental level. Well, in my youth I did some physics that I liked a lot, and the concept of energy in physics is really beautiful. By God, I would love to have invented that concept! Energy has no concrete meaning, like place or time. It is a concept that we constructed ourselves, but it summarises all kinds of things. And then there is conservation of energy. I think in decision theory there is something similar, which I call 'conservation of influence': that's a sort of re-interpretation of preferences in decision making, to make it more natural why all these principles, like the Sure Thing Principle, are good principles. This would involve quite a bit of philosophy: about causation, determinism, and free will. And that should give context to the Bayesian approach. So, you can see, if it is only my intellectual interest, philosophy is close to me and economics is as far as it can be. But my actions are the opposite! (laughs)

CH: Can you say more about the concept of 'influence'?

PW: It governs the decision-making of preference. If I prefer A to B, then my influence is to replace B by A. So I don't say that I prefer A to B, I say that it's my influence that A happens instead of B. And this is the basic start of the concept. It's just a reformulation, but if you use that terminology, all kinds of things change. I have a keyword referring to it in an annotated bibliography on my homepage, so all my fans can follow.

CH: I think I am hardly alone in wanting to see that concept further developed! Would you also relate this to empirical work?

PW: Well, I think that having worked empirically helped me to write in direct and meaningful ways and to notice if I write something that is not clearly verifiable. For me writing a preference axiomatisation and doing an empirical test is about the same. If you know how to measure something, then you

can write preference axioms, and then you also know how to empirically measure them.

CH: Even though there are many axiomatisations that are less well implemented than others...I am thinking of things like the Archimedean axiom.

PW: Indeed, we are in the country of Brouwer, the Dutch mathematician, who influenced me with his constructivism. The Archimedean axiom is not observable, but we sometimes use it because we have the tendency to work with infinite models. And because we work with these models, anytime we work with empirical reality, we are punished by these axioms. We cannot really test continuity. And so we pay a price for this – maybe unfortunate – axiom.

CH: You would like to avoid this?

PW: Yes, definitely! I really like the axiom conditions of finite models, that everything in principle is solvable, and then you discover that getting an EU axiomatisation is much more difficult. The Archimedean axiom makes sense mathematically, but it takes away all kinds of relevant empirical questions. The mathematical question that I would like to solve more than anything else in my life is to have necessary and sufficient preference conditions for expected utility in a finite model. This is very complex, and nobody really knows how to do it. I would like that more than anything else... – So, no Archimedean axiom! (laughs)

CH: That is an excellent closing statement to this fascinating conversation. Thank you very much, Peter!

NEWS

Inferring policy from experiment, 15 May

A mini-conference ‘Inferring policy from experiment’ was held on 15 May at the University of Kent. The event focused on the epistemic and practical issues in using research evidence in policy-design particularly in the areas of medicine and public health, and featured talks by Nancy Cartwright (Durham & UCSD), Sarah Wieten (Durham), and Mike Kelly (Cambridge).

The opening talk of the event was by Nancy Cartwright, titled ‘Two approaches to evidence based health policy—intervention-centred, context-centred’. Both approaches were characterized in terms of their focus questions, target of analysis, and evidential requirements. According to Cartwright’s taxonomy, the intervention-centred approach focuses on characteristics of a policy (Does it work? For whom? What does it cost? etc.), studies repeatable causal processes, and requires evidence in support of causal generalizations. By contrast, the context-centred approach focuses on causal arrangements in the target context of a policy, and studies what causal processes these ar-



rangements afford. This requires robust models of how new intervention-outcome pairs can be brought about given those arrangements. As a consequence the evidential requirements for this approach are demanding.

Cartwright proceeded to argue that the intervention-centred approach makes most sense when the intervention has an ‘in-built’ tendency towards the intended effect. Think for example gravity with respect to the effect of making heavy bodies fall. Whether the intervention-centred approach works for health policy depends on whether health interventions typically have such an inbuilt tendency towards effects that a policy-maker is interested in. For example, there is evidence that deworming programs improve children’s reading scores, but this effect is hardly due to such an inbuilt tendency of the policy-intervention. Rather, the policy has an immediate effect of killing worms, while the effect on reading scores depends on complicated, context-specific causal pathways. The pure intervention-centred approach is thus risky in situations where one lacks knowledge of appropriate supporting causal mechanisms in the target context. It is these supporting causal mechanisms, the analysis of which is the starting point of the context-centred approach. The downside of the context-centred approach is that it may seem prohibitively demanding – learning the details of all the complicated biological, psychological and social mechanisms relevant for the effects of health policy is next to impossible. However, one may learn reliable markers of relevant mechanisms, as well as *cautions* that signal that a given policy might not work in a particular context. Such markers and cautions are not infallible guides to implementation of a policy, but searching for them offers some leverage for dealing with the uncertainty of the pure intervention-centred approach. From this, Cartwright concluded that no matter which approach one adopts, one should hedge one’s bets and plan for failure.

The second talk was by Sarah Wieten. Her talk, titled *What good are pragmatic trials*, offered a critical evaluation of some of the recent arguments in favor of pragmatic trials over explanatory trials. Explanatory trials test an intervention in highly controlled and idealized conditions, thus securing internal validity but arguably compromising external validity as the target population is likely to be dissimilar to the study population. Presumably, pragmatic trials that relax some of the idealized conditions do not suffer from the dissimilarity problem to the same degree. Wieten argued that the arguments in favor of pragmatic trials are, while true in one sense, misplaced with respect to the actual interests of a clinician who will implement the intervention. According to Wieten, the similarity mentioned above is relevant for answering whether the effect of an intervention will be the same in the study and the target populations. But this, according to her, is not the most pressing query from a clinician’s point of view. Rather, one needs an answer to the question: what is the causal effect of treatment in the target population given its observed characteristics. For this query it is important just that the target population does not exhibit features not represented in the study population at all – approaching perfect similarity is not automatically a virtue. This requirement can be achieved in highly idealized trials as well as pragmatic ones. Wieten then elaborated this fundamental point to make two further arguments. Firstly, pragmatic trials seem to deliver the claimed extrapolatory benefits only if the transfer of results from the study population to the target is supported by evidence of underlying mechanisms. Secondly, Wieten argued that the assumed tradeoff between internal and

external validity is not as straightforward as has been claimed, and that there are in fact conditions in which no tradeoff exists at all. In such conditions, choosing a properly idealized study would always be an improvement over the pragmatic approach, as the former will result in higher internal validity.

The final talk of the event was by Mike Kelly, and was titled *Inferring policy from evidence? The case of non-communicable disease and health inequalities in the UK*. Kelly started by reviewing research revealing how geographic and socioeconomic health inequalities in the UK have remained remarkably robust even when the proximate causes of non-communicable disease - such as lifestyle or environmental exposure to pathogens - have greatly varied over time. The question then rises, why has the mounting pool of research evidence on the topic not translated to effective policies for reducing the inequalities? Kelly diagnosed the problem as one of faulty, individualistic epistemology being employed in policy design. When designing policy, health outcomes have typically been conceptualized as properties of atomistic individuals, determined by exposure to risk factors that imping causal effects on individuals in a linear fashion. Crucially, individual choice in light of risks-information has been seen as the most important determinant of exposure to these risk factors. Consequently, health policy regarding non-communicable disease has largely focused on programs for raising risk-awareness in the hope of altering the relevant choice-behavior of individuals for the better – all this in spite of mounting evidence of the role of the wider societal determinants of the distribution of health outcomes.

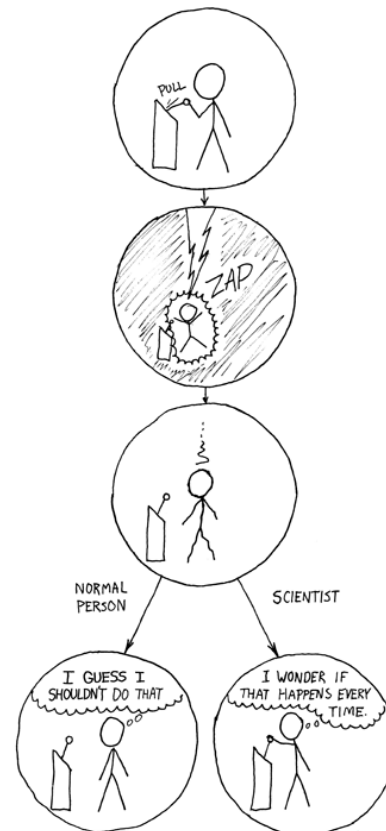
According to Kelly, the individualistic epistemology shifts responsibility to individuals and away from industry, advertisers and the state. It is also mirrored in the methodology of focusing on measuring individual attributes rather than relational ones like social class, gender or social status. As far as social status is considered as a determinant of health, this is again according to Kelly considered as a simple aggregate of individual characteristics, abstracting away from intersecting social attributes such as class or ethnicity. According to Kelly, what is needed to improve the situation is a move towards dynamic thinking that involves explicit analysis of history, power, gender and class relations, social justice, as well as the biological mechanisms through which the health effects of socioeconomic differences are transmitted and maintained across generations.

VELI-PEKKA PARKKINEN
Philosophy, Kent

Perspectives on Explanation, 18–19 May

“Perspectives on Explanation” was an international workshop held in Prague on May 18–19, 2017, organized by the Department of Analytic Philosophy of the Institute of Philosophy of the Czech Academy of Sciences, and supported by the grant project Formal Epistemology: Future Synthesis. The idea of organizing this workshop as an event, which has its roots in the newly formed East European Network for Philosophy of Science, was conceived by Daniel Kostic and Martin Zach.

Ladislav Kvasz (Czech Academy of Sciences) gave the opening talk. He began with two examples of explanation taken from the history of mathematics, showing how something from one particular (mathematical) viewpoint (e.g., geometric) once seemed incomprehensible but later found a natural explanation within a newly constructed viewpoint (e.g., algebraic). Expla-



nation, as Kvasz argued, is a linguistic phenomenon which obtains between two linguistic frameworks. Kvasz then introduced the notion of explanatory power, which, according to him, shows how the language allows us to explain the failures which occurred in the previous stages, failures that were previously incomprehensible. This was then illustrated on a number of examples taken from the history of physics.

Second speaker, Sorin Bangu (University of Bergen), first introduced the example of superconductivity and sketched the history of attempts at explaining this phenomenon (i.e., London brothers model, Ginsburg-Landau model, BCS model). Based on that, Bangu argued the case for higher-level explanations in physics in which fundamental laws and entities play the role of causal background but are not genuinely explanatory relevant. Bangu thus argued against the idea that explanation is transitive, at least not always as was illustrated on the superconductivity case. Simply put, one cannot derive an explanation of superconductivity solely from the fundamental level, but rather, one has to focus on the explanatory relevant level.

Robert Batterman (University of Pittsburgh) opened the afternoon session with a talk on universality, stability, autonomy, and scales. In his talk, Batterman defended the view that explaining certain phenomena (e.g., the behavior of critical phenomena) requires recognizing that there is a certain universal behavior displayed by multiple systems that can and often are completely different from each other on the micro-level. Furthermore, this universality is stable under perturbation. Batterman attacked what he calls the ‘common features account of explanation’ and answered some of the critiques directed at his own view, the so-called minimal model explanation. According to Batterman, although the common features might be necessary for explaining the universal behavior, they are not sufficient; one has to explain why these features and not some

others are adequate for the explanation and this is done by actually demonstrating the irrelevance of a number of micro-details.

Arnon Levy (Hebrew University of Jerusalem) distinguished between pragmatic and what he calls ‘factualist’ account of explanation. Levy argued that a proper account of explanation combines elements from both the pragmatic and factualist views and that, in fact, both derive from understanding where understanding is the ability to make counterfactual inferences (which, in turn, makes understanding non-subjective). Understanding, as Levy claimed, is an ability to make inferences. Levy then further argued that his account provides a rationale for explanatory idealizations, something that the traditional pragmatic views dismiss, as well as for causal modularity.

Alisa Bokulich (Boston University) reviewed the ontic conception of explanation and showed that the proponents are inconsistent in their attempts at answering objections to their view. Bokulich then argued for what she calls an ‘eikonic conception of explanation’. She presented four examples taken from different scientific fields which motivated her conclusion that the role of representation, which pertains to both explanans and explanandum, is indispensable in the explanatory practices of scientists. Based on the examples, Bokulich then further argued that it is often the case that no ‘one best’ representation allows scientists to answer every question about a given entity. She closed by showing a number of benefits of the eikonic conception of explanation.

Marcin Milkowski (Polish Academy of Sciences) noted that there does not seem to be any explanatory unity in cognitive science. Instead, we have different ‘paradigms’, though these have little to do with the original Kuhnian sense since incommensurability is not even suggested. Milkowski then considered a number of such ‘paradigms’ or approaches (e.g., embodied cognition). These approaches are not like grand unified theories; they are akin to methodological principles. They are very broad but they also rely on empirical evidence, which, according to Milkowski, makes them something like research programs in Lakatos’ sense. In conclusion, Milkowski denied that either integration or unification could bring a primordial unity to cognitive science. ‘Paradigms’ are simply taken to be certain general hypotheses that are supposed to drive future research.

Lilia Gurova (New Bulgarian University) argued for what she calls the ‘inferentialist view of understanding’. Such a view consists of defending the claims that good explanations increase our understanding of the phenomena and that understanding is best analyzed in terms of inferences one can draw about the phenomena. Gurova identified problems for both factivist and non-factivists accounts (subjective as well as objective versions) of understanding. The inferentialist account defended by Gurova agrees with non-factivists that false theories/models may provide understanding but parts way on the question of relevance of truth; for Gurova, truth is not irrelevant since it may allow for extra-inferences.

[Here](#) you can find the recordings of all the talks. Some of the photos from the workshop can be found [here](#).

MARTIN ZACH

Charles University in Prague

Calls for Papers

FORMAL AND TRADITIONAL EPISTEMOLOGY: special issue of *MANUSCRITO*, deadline 1 July 2017.

LOGIC, INFERENCE, PROBABILITY AND PARADOX: special issue of *Philosophies*, deadline 20 July 2017.

NEW TRENDS IN RATIONAL CHOICE THEORY: special issue of *Topoi*, deadline 27 August.

FOUNDATIONS OF CLINICAL REASONING: AN EPISTEMOLOGICAL STANCE: special issue of *Topoi*, deadline 31 August.

KNOWLEDGE AND JUSTIFICATION: NEW PERSPECTIVES: special issue of *Synthese*, deadline 1 September.

REASON & RHETORIC IN THE TIME OF ALTERNATIVE FACTS: special issue of *Informal Logic*, deadline 1 September.

WHAT IS A COMPUTER?: special issue of *Minds and Machines*, deadline 30 September.

THE REASONER SPECULATES

RWT in the social sciences

We speculate that the Russo-Williamson thesis (RWT) can be extended to the social sciences and has some interesting consequences there, both for evidence-based policy and for the debate between advocates of qualitative and quantitative methods.

Russo and Williamson (2007: [Interpreting causality in the health sciences](#), *International Studies in the Philosophy of Science* 21(2), 157–170) put forward the idea that in the health sciences, in order to establish that A is cause of B one needs to establish both that A and B are appropriately correlated and that there is some underlying mechanism linking A and B that can explain instances of B in terms of instances of A and that can account for the correlation between A and B .

Russo and Williamson (2007) argued that this thesis has consequences for the metaphysics of causality: standard difference-making accounts of causality, as well as mechanistic theories and pluralist accounts all struggle to explain RWT. RWT also has consequences for the health sciences. Clarke et al. (2014: [Mechanisms and the Evidence Hierarchy](#), *Topoi* 33(2), 339–360) used RWT to argue that evidence-based medicine mistakenly takes clinical studies to be evidentially superior to other evidence of mechanisms. If RWT is correct, evidence of mechanisms arising from sources other than clinical studies should be treated on a par with clinical studies. This observation led to the [EBM+](#) programme.

We speculate that RWT also has interesting consequences for the social sciences. This is because RWT applies more widely than the health sciences. It is true in general—not just in the health sciences—that establishing correlation is insufficient for establishing causation: one needs to establish that the correlation is explained by a genuine causal connection rather than by, say, bias, confounding or some other non-causal link between the two variables. The existence of a mechanism of action is characteristic of a genuine causal connection, so establishing



causation requires establishing the existence of a mechanism as well as a correlation.

RWT has two sorts of consequence for the social sciences. First, the lessons concerning evidence-based medicine also apply to evidence-based methods in the social sciences. In evidence-based public policy, for instance, statistical studies which measure instances of the putative cause and effect A and B —especially randomised controlled trials (RCTs)—tend to be viewed as strictly superior to other kinds of evidence. (See, e.g., the UK [What Works Network](#).) However, while these studies do indeed provide good evidence for the existence and the extent of a correlation between A and B , they typically do not suffice to establish causality. Other kinds of evidence of the underlying mechanisms is usually also critical.

For example, many forms of bias affect RCTs in the social sciences. One of the most discussed forms of bias is the Hawthorne effect: when a trial is not perfectly double-blind, members of the experimental group may be aware of being observed. Cilliers et al. (2015: [The White-Man Effect: How Foreigner Presence Affects Behavior in Experiments](#), *Journal of Economic Behavior & Organization* 118, 397–414) illustrates how this may lead individuals to change their behavior for the duration of the trial. In such a situation, mechanistic evidence can reveal that an outcome B is caused not only by the treatment A , but also by the reaction to the experimental setting.

The second consequence of RWT for the social sciences concerns the debate between advocates of qualitative methods and advocates of quantitative methods. In the last few decades, there has been an enormous amount of discussion of the relative virtues of quantitative and qualitative methods. While some authors have claimed that such methods are incompatible, recently many researchers have argued that quantitative and qualitative methods are best thought of as complementary and should therefore be mixed in order to combine their strengths.

RWT provides a new angle from which to examine this controversy. If RWT is correct, the mixture of methods that are required to establish causality is precisely the mixture required to establish the existence of a correlation and a mechanism. In some cases, large-scale statistical studies—i.e., quantitative methods—suffice to establish causality. In particular, when the studies are sufficiently large, well constructed and well conducted, and independent studies all find a correlation, and this correlation is sufficiently large, then bias, confounding and non-causal explanations of the correlation can sometimes be ruled out. Typically, however, the available studies that measure both A and B do not suffice to establish causation. In these cases, one needs other evidence, including both qualitative and quantitative evidence of the underlying mechanism.

For instance, Weller and Barnes (2016: [Pathway Analysis and the Search for Causal Mechanisms](#), *Sociological Methods & Research* 45(3), 424–457) point to limitations associated with evidence of correlation and defend the combined use



of quantitative and qualitative methods to establish causation. Starting from the statistical relationship between natural resources and civil conflicts, they explore case studies to investigate the presence of underlying mechanisms.

So, should one adopt qualitative, quantitative or mixed methods when establishing causality in the social sciences? If RWT is correct, this depends. In some cases, quantitative methods suffice, as we have just seen. In other cases, qualitative methods may even suffice. Although this avenue is less clear, as a possible example consider Schimmelfennig (2001: [The community trap: Liberal norms, rhetorical action, and the eastern enlargement of the European Union](#), *International organization* 55(1), 47–80), who used only qualitative methods to detect three mechanisms that, together, caused some countries which were initially against eastern enlargement of the EU to eventually support it. Typically, however, one needs a mixture of methods and the methods required are those that establish both correlation and mechanism. RWT thus serves to shed light on the controversy between qualitative, quantitative and mixed methods in the social sciences.

VIRGINIA GHIARA
JON WILLIAMSON

Philosophy, University of Kent

REVIEWS

(Formal) Argumentation Theory

As a regular reader of the Reasoner, it was with some trepidation and enthusiasm that I accepted Hykel's invitation to contribute a regular column on 'what's hot in argumentation'; trepidation at the prospect of meeting the high standards set by current and past contributors, while enthusiastically seizing the opportunity to report from the front line of a field whose multidisciplinary reach and potential has never been more relevant in a post truth world of echo chambers and concerns about the dangers of AI (of which more in this and later columns). In this column I set the scene for future dispatches, and while the focus here is on 'formal' logic-based models of argumentation, it is the very appeal of these models that they aim to straddle the divide between normative, logic-based models of reasoning and more descriptive accounts of the way humans reason; hence the title's inclusion of 'formal' in parentheses.

In its classical treatment within philosophy, from Aristotle to the present day, the study of argumentation has focussed on what with some poetic license I refer to as 'human orientated' uses of argument, such as when an argument can be considered legitimate or flawed, the processes by which participants engage in debate, rhetorical aspects of argumentation e.t.c. However, in the last two decades, the study of logic, and more broadly 'computational', models of argumentation has emerged as a growing sub-area of AI, as witnessed by dedicated journals, conferences, workshops and national and international projects on argumentation. While researchers from a number of fields, including philosophy and legal reasoning, contributed to this growth, it was primarily those with a background in non-monotonic reasoning that fuelled this surge of interest in argumentation.

AI research in the 80s and early 90s saw a proliferation of non-monotonic logics tackle classical logic's failure to formalise our common-sense ability to reason in the presence

of incomplete and uncertain information. In the classical paradigm, the inferences from a set of formulae grows monotonically, as the set of formulae grow. However in practice, conclusions that we previously obtain may be withdrawn because new information conflicts with what we concluded previously or with the assumptions made in drawing previous conclusions. Essentially then, a key concern of non-monotonic reasoning is how to arbitrate amongst conflicting information; a concern that is central to the argumentative enterprise. It is this insight that is substantiated by argumentative characterisations of non-monotonic inference. Most notably Dung, in his seminal paper (P. M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.), considered arguments constructed from a given set of formulae Δ . An argument consists of formulae in Δ supporting the argument's claim (essentially each argument is a self-contained proof of a conclusion derived from the supporting formulae). These arguments are then related to each other in a directed graph (Dung framework), based on whether one argument is a counter-argument (*attacks*) another; for example when the claim of one argument negates a formula in the support of the attacked argument. In this way the formulae Δ are said to 'instantiate' the framework. Sets of acceptable arguments (*extensions*) are then identified under different 'semantics', where the fundamental principle of defense licenses membership of an argument in any such extension: an argument X is in an extension E if every argument Y that attacks X is itself counter-attacked (i.e., defended) by some $Z \in E$. The different semantics capture additional criteria for membership, for example whether the defending Z is or is not required to be distinct from X . Arguments in the extensions are thus justified, and their claims identify the non-monotonic inferences from Δ . Dung and others show that the inferences defined in this way correspond to the inference relations of various non-monotonic logics defined directly over the given sets of formulae Δ .



A vast body of research builds on Dung's theory, in part because the foundational principle of defense that underlies these dialectical characterisations of non-monotonic inference is compatible with everyday intuitions about how we reason and debate. Thus the theory has been extended to incorporate features of commonsense reasoning neglected by the non-monotonic logic community, via the modelling of features of human orientated argumentation. For example: the use of *schemes* and *critical questions* – stereotypical patterns of argument and counter-argument developed by the philosophical community – in guiding argument construction and identification of attacks; the modelling of support relations amongst arguments; evaluating arguments by additionally accounting for the cumulative weight of multiple arguments in support of a claim, or preferences reflecting the relative trustworthiness of sources of arguments, or the relative importance of values promoted by the actions warranted by arguments in practical reasoning.

Moreover, I argue that these dialectical characterisations of non-monotonic inference provide for a more compelling ac-

count of the relationship between proofs (syntax) and models (semantics). Given a Dung framework, argument game proof theories establish whether a given argument X is justified. The essential idea is that a proponent wins a game iff she successfully counter-attacks (defends) against all attacking arguments moved by an opponent, where all attacks moved are licensed by reference to those represented in the given framework. A game is won in respect of showing that X is justified, iff X is justified in the sense that it belongs to an extension of the framework under some semantics (with rules on the allowable moves in the game varying according to the semantics). Contrast this with the relation between syntax and semantics traditionally adopted by non-monotonic logics, that by virtue of being rooted in and extending the classical paradigm, are concerned with showing the correctness of inferences delivered by proof theories from given sets of formulae, in the sense that the inferences are 'true' in models of these formulae rendered in distinct mathematical structures (e.g., interpretations, possible worlds, algebraic structures). However, the argumentative paradigm can be seen to blur the distinction between syntax and semantics, and this, I argue, is as it should be. For the distinction and sought for correspondence between proofs and models is a legacy of the epistemological distinction between belief and knowledge, whereby our beliefs are said to meet a standard of correctness to the extent that they correspond with knowledge, interpreted as justified true belief. Of course, the assumption of access to some objective characterisation of what is true, when establishing such correspondences, is to say the very least highly problematic. However, as an advocate of Popper's epistemology, I find that a more pragmatic and coherent proposition is that what is believed true is *provisionally* true to the extent that all attempts to *thus far* show otherwise have failed; no access to some objective standard of truth need be assumed. In argumentation terms – specifically argument game proof theories – the claim of a justified argument is provisionally true to the extent that all counter-arguments thus far considered have been defended (and all defenders in turn are defended, and so on); no appeal need be made to some mathematically distinct model theoretic structure. Indeed, we might take more seriously the idea that 'it is reasoning (qua proof theory) all the way down', and abandon reference to a *given* framework of arguments when establishing the justified status of an argument in the above described two player games. Rather, it is the arguments moved that *incrementally define a framework*. More precisely, the contents of moved arguments – the contained propositions – incrementally define the set Δ of formulae that instantiate a framework. This idea would admit a somewhat radical (although in my view intuitive) interpretation in the context of an individual reasoner considering arguments for and against some claim: the contents of her beliefs reveal themselves as they are deployed in arguing the matter. What is less contentious is the generalisation of the idea to another important feature of commonsense reasoning neglected by non-monotonic logics; that no man is a (reasoning) island.

We reason through dialogue, debate and discussion, for example when persuading and deliberating with others as to the cogency of beliefs or appropriateness of actions, and formal models of argumentation naturally generalise to dialogical models of distributed non-monotonic reasoning. These dialogical models can be seen as generalising argument games. Interlocutors exchange locutions that explicitly consist of arguments, or implicitly define arguments (for example, when

making a claim and then when questioned as to its veracity, providing supporting reasons). The contents of these locutions incrementally instantiate a framework, such that one then aims for a correspondence between evaluating the outcome of a dialogue in favour of some claim, just in case the claim is supported by some justified argument in the framework. In other words, when the claim is a non-monotonic inference of the locutions' content. There is much potential for development of these models to support not only machine dialogue, but human, and human-machine dialogue, and realising this potential will require further interdisciplinary collaboration with (amongst others) researchers in natural language processing, speech act theory and human orientated approaches such as the pragma-dialectic school of argumentation. Indeed, I suggest that the need for supporting human machine dialogue has never been more pressing, given the widely publicised concerns about the dangers of AI, and what has been termed the 'value loading/alignment problem'.

The idea that future AIs may pose a threat to human well being, by virtue of unforeseen consequences of actions purposed to achieve their operators' goals, is not new (this is after all the rhetorical thrust of Asimov's three laws of robotics stories). However, the issue has acquired renewed urgency given the startlingly successful use of machine learning techniques in recent years. A *feature* of learning systems is the discovery of unforeseen ways of achieving goals, and as argued by Nick Bostrom in *Superintelligence*, the single minded achievement of *any* operator's goal will incentivise AIs to thwart corrective measures to prevent harm. To address this problem, current efforts in the machine learning community focus on machines effectively learning human values and preferences through observation and query. This I argue will not suffice. For the assumption is that observed actions reveal preferences and hence values, and that humans are sufficiently informed and have the requisite capacities to definitively arbitrate on matters of ethical importance. However humans clearly do not always behave ethically, and moreover are often uncertain about how to resolve ethical issues; particularly those arising from deployment of novel technologies (that hence lack precedent). What is needed is many minds, both artificial and human, to ensure alignment of machine actions with human values. We require that AI systems and humans engage in comprehensive, rational exchange of arguments purposed to decide ethical issues; indeed, *in dialogues that will be better purposed to do so* by virtue of incentivising and harnessing AI's vastly superior access to information and capacity to look further into the future, while accounting for human reasoning about values.

SANJAY MODGIL


Informatics, King's College London

WHAT'S HOT IN . . .

Uncertain Reasoning

Why did Alice go to the fridge? Because she *believed* there was a beer in the fridge, and she *desired* a beer. This seems like a perfectly good explanation of her behaviour. Why did Bob bet on *Categorical Imperative* to win the horse race at 4 to 1? Because Bob's degree of belief that *Categorical Imperative* would win was greater than 0.2. This, too, seems like a reasonable account of a rational agent's behaviour. But the relationship

between these two theories – full beliefs and partial beliefs – been a source of philosophical puzzles for quite some time now. Leitgeb's new book *The Stability of Belief* (OUP, 2017) engages with this question. The core problem of the book is how to reconcile these two pictures of rational belief. Leitgeb's solution – his stability theory of belief – is a clever and mathematically sophisticated approach that I won't go into here. I want to pick up on something Leitgeb says near the beginning of the book.



Leitgeb canvasses some positions you might take towards the relationship between full belief and partial belief in Chapter 1. One option is to assert that only one kind of belief refers: there is no such thing as full belief or there is no such thing as partial belief. That both explanations I started the last paragraph with seem acceptable tells against this "eliminativist" position. The second option Leitgeb discusses is that both full belief and partial belief refer to something, but one does so derivatively, and can be reduced to the other. Maybe full beliefs are just sufficiently high degrees of belief, or maybe degrees of belief are just full beliefs about chances. Neither of these suggestions really pans out, but perhaps some more sophisticated reduction could tidy up some conceptual clutter. The third option Leitgeb discusses is that both full belief and partial belief refer, and they refer to different things. So we have a full belief system and we have a partial belief system in our brains, and they function (somewhat) independently of each other.

Reading this overview of positions one might take towards the relationship between full and partial belief reminded me of the story about the blind men and the elephant. Each blind man is directed towards one part of the elephant – tusks, leg, trunk, ears... – and each one gets a very different view about what sort of thing an elephant is. In forming their opinion about what an elephant is from their limited understanding of it, are these characters referring to the same thing or to different things? To me, this seems like a false dichotomy: they are referring to or theorising about different parts of one thing. You can't "reduce" tusk-elephant theory to ear-elephant theory, and yet they aren't independent entities either. I think the same is true of the theories of belief that Leitgeb discusses.

Full belief and partial belief are both incomplete, approximate, imperfect representations of part of a complex and messy mental state. They are not representations of the same thing, but they are not representations of different things either. I don't take myself to be disagreeing with Leitgeb in making this point. Even in discussing the possibility that full belief and partial belief are "irreducible", Leitgeb accepts that there must be some connections between the two. If the full belief and partial belief are just two distinct compartments in the mind, then we have a problem of how and why the two connect or should connect. But if we think of full and partial belief as two imperfect representations of different aspects of the same phenomenon, then the covariation in the two can be explained by them both representing the same thing. When the elephant legs start moving North, the elephant ears also move North, and by the same distance. This isn't a mysterious covariation that needs explaining,

but if we took the blind men's theories partial elephant theories to be independent, then it would appear mysterious.

I don't think Leitgeb would necessarily disagree with anything I've said here. The official position of his book is agnosticism between reduction and irreducibility. The position I've suggested here is something of a happy medium between those extremes.

SEAMUS BRADLEY

Philosophy, University of Tilburg

Philosophy and Economics

In this column, I will regularly review some trends in the flourishing field of interdisciplinary research in and between philosophy and economics. Research in this area really has increased recently, as has general interest in philosophical questions about economics, as well as economic methods being employed to answer philosophical questions. There are now a number of journals, regular conferences, societies, and networks that all cover different aspects of research between philosophy and economics. Rather than attempting a general overview or exploration, I will simply pick up different topics and themes as I go along.

This time around, I will start 'at home', as it were, and focus on some goings on at the Erasmus Institute for Philosophy and Economics (EIPE) at the Faculty of Philosophy at Erasmus University Rotterdam. EIPE was founded in 1997 and is currently celebrating its 20th anniversary. To kick off the anniversary celebrations, we teamed up with the local Erasmus Happiness Economics Research Organization (EHERO) to organise a week-long interdisciplinary conference in March 2017 that focused on the theme of happiness and wellbeing. Two of the keynote speakers, Anna Alexandrova (Cambridge) and Erik Angner (Stockholm), gave stimulating talks on the conference theme. Erik focused on facilitating the dialogue between empirical happiness researchers on the one hand and philosophers of science on the other hand. Anna inquired about the role of expertise with regards to well being, criticising methodological choices during validation of measures in happiness research. In seven parallel sessions, more than 30 talks were given on conceptual issues, measurement, decision theory, inequality, and welfare economics in connection to the wellbeing and happiness theme. We also witnessed some very interesting exchanges between philosophers and economists working on happiness and wellbeing.

To talk about issues in happiness and wellbeing was not the only aim of the conference. When EIPE was founded in 1997, by a group of philosophers and economists at Erasmus, one of the aims was to provide training for the next generations of interdisciplinary researchers. Accordingly, an MPhil and PhD programme in philosophy and economics marked the beginning of the institute. From both programmes, there are now a healthy number of alumni, many of whom are working in philosophy and economics, and many of them also attended the conference. Indeed, two further invited speakers, Johanna Thoma (London School of Economics) and Francois Claveau (University of Sherbrooke), both Alumni of EIPE, gave keynote talks and commented on each other. The format of their keynote session, with quick succession of (i) talk, (ii) comments, and (iii) discussion with the audience was similar to the EIPE seminars, in which senior researchers and PhD students alike are usually commenting on each others' papers in this format. First up

was Johanna with a talk on transitivity on which Francois commented, with discussion thereafter. Next, Francois talked about expertise, and Johanna commented, followed by discussion. It was a welcome variation of format for a keynote session, and it also required a lot of philosophical skill from both protagonists!

The conference also featured contributed papers in a standard parallel sessions format. In total, the programme contained more than 70 contributed talks in over 16 sessions on all aspects of the philosophy of economics. On the ethics side, there were several sessions on distributive justice and the capability approach, in addition to the sessions relating to happiness and wellbeing. On the decision theory side, five different sessions examined the foundations of rational choice theory. On the methodology side, six sessions focused on explanation, mechanisms, measurement, evidence, and causality. And two sessions investigated the role of philosophers and economists in each other's disciplines. All together, these talks showed just how far the discipline of philosophy and economics has become, and how diverse and rich it now is, with plenty of upcoming younger scholars. I will certainly return to some themes in future columns!

EIPE is also home to the Erasmus Journal for Philosophy and Economics (EJPE). EJPE is a peer-reviewed open access journal that publishes high quality original research in all areas of philosophy of economics, as well as interviews with leading scholars, and summaries of recent PhD theses in the area of philosophy and economics. The journal also runs an annual competition for the Mark Blaug Prize for younger scholars. The journal is edited by doctoral students of EIPE, and will publish a special issue with papers from the conference. You can find its [website](https://ejpe.org/) at <https://ejpe.org/>. I can only encourage you to check out its articles and interviews.

From the next issue of THE REASONER onwards, I will select some specific topics and trends in philosophy and economics and review them here. So long!

CONRAD HEILMANN

Erasmus Institute for Philosophy and Economics (EIPE)

Erasmus University Rotterdam

Evidence-Based Medicine

Last year, a [systematic review](#) was published in *PLOS ONE*. The aim of the review was 'to evaluate the effects of school-based educational interventions for enhancing adolescents' abilities in critically appraising health claims'. However, the authors concluded that 'serious limitations in the existing evidence make it difficult to draw definitive conclusions concerning the effect of school-based educational interventions for enhancing adolescents' abilities to critically appraise health claims'. In particular, the authors of the study pointed out the following limitations: A small number of studies was included in the review, and there was a high risk of bias in these included studies. In addition, the authors pointed out that '[n]one of the studies measured students' appraisal behaviour in everyday contexts outside the classroom, which would be the ultimate goal of improving students' abilities to critically appraise health claims in society'.


The authors proposed the following implications for future research:

The results of this systematic review indicate that there is a lack of school-based educational inter-

ventions for enhancing critical appraisal abilities of health claims among adolescents. Thus, novel interventions that aim to improve and sustain these abilities should be developed and evaluated. Well-designed evaluation studies are needed; preferably pragmatic cluster-randomised controlled trials that take place in a wider variety of school-based settings and that closely resemble normal educational practice.

Recently, some researchers have attempted to address this need. Last month, a [cluster-randomised controlled trial](#) was reported in *The Lancet*. The trial aimed ‘to evaluate an intervention designed to teach primary school children to assess claims about the effects of treatments’. A random selection of schools was assigned the intervention of a number of lessons during a single term, where the teachers received some relevant training and were provided with a set of teaching resources including textbooks. The control group of schools received no intervention. The schools were then given a multiple-choice test. And the intervention schools performed significantly better in the test than the control schools. The authors concluded that the intervention leads to an improvement in children’s assessment of health claims. This is clearly a step in the right direction. Although there is still a concern about whether outside of the classroom there is also an improvement in the children’s assessment of health claims.

MICHAEL WILDE
Philosophy, Kent



IF YOU KEEP SAYING "BEAR WITH ME FOR A MOMENT", PEOPLE TAKE A WHILE TO FIGURE OUT THAT YOU'RE JUST SHOWING THEM RANDOM SLIDES.

EVENTS

JUNE

- [PRFS](#): International Workshop: Proofs, Paris, 1–2 June.
- [P&JB](#): Perception and Justified Belief, Ruhr-University Bochum, Germany, 1–2 June.
- [VR-VA](#): Visual Reasoning – Visual Arguments, Greifswald, Germany, 1–3 June.
- [PoP](#): Philosophy of Probability Graduate Conference, London School of Economics, 2–3 June.

- [ISW](#): Imagination in Science Workshop, University of Leeds, 6 June.
- [WAAE](#): Wittgenstein and Applied Epistemology, Lisbon, 6–7 June.
- [TACrS](#): Time and Causality in the Sciences, Stevens Institute of Technology, 7–9 June.
- [STE](#): Simulation and Thought Experiment, University of Geneva, 8–9 June.
- [CAP](#): Conditionals and Probability, University of Leeds, 10–11 June.
- [OoER](#): The Ontology of Epistemic Reasons, University of Basel, 14 June.
- [PIS&S](#): Progress in Science and Society, Workshop with Philip Kitcher, Leibniz University Hannover, 14 June.
- [NoL](#): The Normativity of Logic, University of Bergen, Norway, 14–15 June.
- [E&DMIL](#): Evidence and Decision Making in the Law, King’s College London, 16 June.
- [FW&AtDO](#): Free Will and the Ability to Do Otherwise, Campus Belval, Esch-Belval, Luxembourg, 16–17 June.
- [LEARNAUT](#): Learning and Automata, Reykjavik, Iceland, 19 June.
- [CAM](#): Conceivability and Modality, Sapienza University, Rome, 19–20 June.
- [CEC](#): Causation, Explanation, Conditionals, LMU Munich, 21–23 June.
- [EUADM](#): Evidence, Uncertainty and Decision Making with a Particular Emphasis on Climate Science, University of Salzburg, 22–23 June.
- [GT&DT](#): Workshop on Algorithmic Game Theory and Data Science, Cambridge, Massachusetts, 26 June.
- [CCC](#): Continuity, Computability, Constructivity—From Logic to Algorithms, Nancy, France, 26–30 June.
- [SoML](#): 17th Latin American Symposium on Mathematical Logic, The Benemérita Universidad Autónoma de Puebla, 26–30 June.
- [EoM](#): Epistemology of Metaphysics, Prague, 29–30 June.
- [UT](#): Underlying Thought: Philosophical Analyses of Epistemic and Ethical Cognition, Cardiff University, 29–30 June.
- [LCiCT](#): London Conference in Critical Thought, London South Bank University, 30 June–1 July.

JULY

- [VoKH](#): The Varieties of Knowing How, Essen, Germany, 6–7 July.
- [AAFL](#): Conference of the Australasian Association for Logic, University of Adelaide, 6–8 July.
- [PA](#): Philosophical Analysis, Krakow, Poland, 7–9 July.
- [PLS](#): Panhellenic Logic Symposium, Delphi, Greece, 12–16 July.
- [OCL&TM](#): Workshop in OCL and Textual Modeling, Marburg, Germany, 20 July.

AUGUST

- [MLwG](#): Mining and Learning with Graphs, Halifax, Nova Scotia, Canada, 14 August.
- [CW](#): Causality Workshop: Learning, Inference, and Decision-Making, Sydney, Australia, 15 August.
- [LFoUAL](#): Logical Foundations for Uncertainty and Learning, Melbourne, Australia, 19 August.

PLAAM: Philosophy, Logic and Analytical Metaphysics, Brazil, 21–23 August.

COURSES AND PROGRAMMES

Courses

COMPUTER SIMULATION METHODS: Summer School, High Performance Computing Center Stuttgart (HLRS), 25–29 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 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: 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 METHODS AND PRACTICES OF PHILOSOPHICAL RESEARCH: Northern Institute of Philosophy, University of Aberdeen.

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

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

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

MA IN REASONING

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.

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.

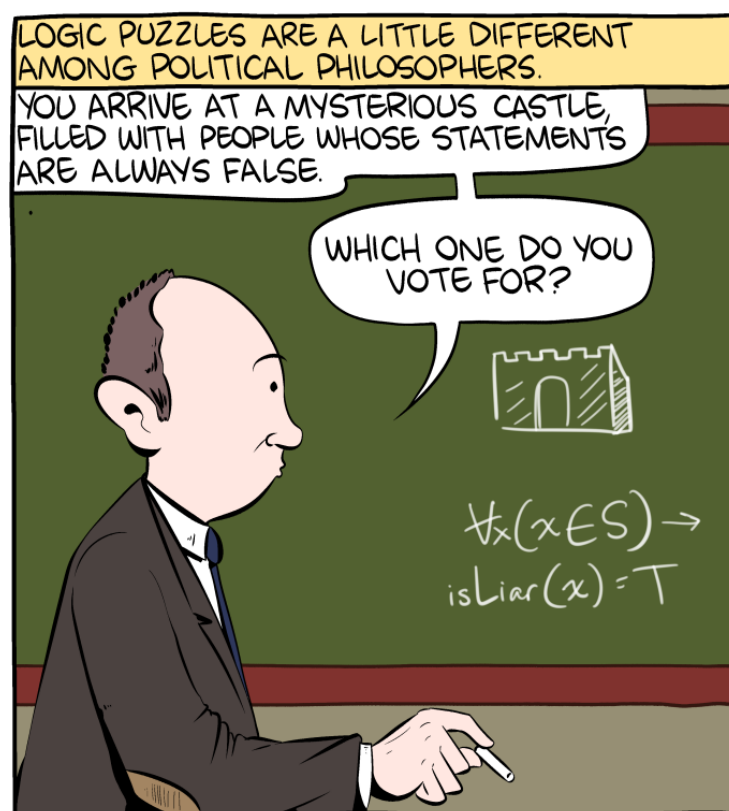
MSc IN LOGIC: Institute for Logic, Language and Computation, University of Amsterdam.

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.



JOBS AND STUDENTSHIPS

Jobs

SENIOR LECTURER: in Clinical Trial Statistics, University of Manchester, deadline 4 June.

RESEARCH ASSOCIATE: in Machine Learning and Neuroimaging, University College London, deadline 8 July.

PROFESSORSHIP: in Statistics/Probability, the University of Copenhagen, deadline 11 June.

LECTURESHIP: in Philosophy of Mind and/or Epistemology, the University of Bristol, deadline 25 June.

Studentships

PHD: in Data Analytics and Society, University of Leeds, deadline 6 June.

PHD: in Statistics, University of Southampton, deadline *open*.

PHD: in Epistemology/Philosophy of Mind, University of Fribourg, Switzerland, deadline 30 September.