The Reasoner

Volume 11, Number 9
September 2017

thereasoner.org
ISSN 1757-0522

Contents

Editorial

It is with enormous pleasure that I accept Hykel's request to serve as guest editor for this issue of The Reasoner, and get the chance to interview Rineke Verbrugge. Rineke holds the chair of Logic and Cognition at the Institute for Artificial Intelligence and Cognitive Engineering (ALICE), Faculty of Science and Engineering of the University of Groningen. Interdisciplinarity is the 'bread and butter' of the Reasoner's readership and Rineke's research is no doubt a beautiful example of the sort of original insights interdisciplinarity can lead to in the study of reasoning and decision-making. Rineke has been a true pioneer in crossing seemingly insurmountable disciplinary borders, such as those separating Logic and Game Theory on one side, and Cognitive Sciences on the other. This type of attitude running against disciplinary silos is a central thread in the interview I am sharing here. I hope you enjoy learning about Rineke's views as much as I did.

Davide Grossi
Computer Science, University of Liverpool

Features

Interview with Rineke Verbrugge

DAVIDE GROSSI First of all thanks for accepting to be interviewed for the Reasoner. Let me start with an introductory question. As scientists we are used to label ourselves all the time. We are used to say things like, "I'm a ...-ist" (computer scientist, cognitive scientist, etc.) or "a ...-ian" (mathematician, logician, etc.) or "a ...-er (e.g., philosopher, etc.). It seems hard to make a choice in your case. Do you use such labels at all? And can you tell our readers a bit about your background?

RINEKE VERBRUGGE I don't like labels that much, and I prefer not to introduce myself as "I am an x", no matter what the x may be. But I do sometimes refer to myself as "a logician by training". What I by far prefer to do is to rather focus on what my current interests are and to say things like "I am currently interested in the interface between logic and cognitive science" or "I am currently working on an experiment based on game theory and computational cognitive models". Labels often lead to essentialist generalizations like "logicians are X so computer scientists must be not X" with all the disservice this causes to research communities. Definitely such labels would lead to a lot of contradictions in my case. As to my background, I liked mathematics and logic since I was 14 or 15

years old, when I was living for one year in Houston, US. My math teacher there made me enthusiastic for logic. When back in the Netherlands I read Douglas Hofstadter's book Gödel, Escher, Bach, which just came out around that time. I remember that Gödel's incompleteness theorem really fascinated me. I also recall that when I was 16, a special issue of the Dutch magazine Vrij Nederland came out on the Dutch mathematician Luitzen Egbertus Jan Brouwer, who invented intuitionism and was also a great topologist. The special issue consisted of eminent mathematicians of that time who made their best attempts to explain the mathematics, logic and philosophy behind Brouwer's work. That also completely fascinated me. So at that time I decided I would go and study at the University of Amsterdam, where Brouwer's tradition was strongest, and where the scientific grandson of Brouwer, Anne Troelstra, was working at the time. I thoroughly enjoyed my mathematics studies there, topology, analysis but especially logic and foundations of mathematics. I also took side courses in philosophy, mostly philosophy of language, Wittgenstein, Montague grammar, and modal logic.

DG The foundations of mathematics were also the topic of your PhD research, right?

RV Indeed, and also of my master thesis. I did both of them mostly with Dick de Jongh and Albert Visser, who were my main supervisors at that time. Anne Troelstra was my promotor, looking so to say, from the sidelines. At that time there was a nice group on the foundations of arithmetics in Amsterdam, looking at provability logic, interpretability logic, including relations with computational complexity theory.

DG Was that at the institute that Evert Beth founded in Amsterdam and that then became the Institute for Logic, Language and Computation (ILLC)?

RV Yes that was becoming the ILLC around that time in the 90s. It was the "Instituut voor Grondslagenonderzoek en Filosofie der Exacte Wetenschappen" (Institute for the research on the foundations of mathematics and philosophy of the exact sciences). Since Johan van Benthem came to the institute, around 1986, he has been instrumental in expanding the remits of that institute towards a more interdisciplinary profile. It started as ITLI "Instituut voor Taal Logica en Informatie" (Language, Logic and Information) to later become the ILLC. Those were very exciting times! And even though I later changed my research focus, foundational questions have remained a keen interest for me, and still occupy some of my free time. For instance, I followed with much interest the latest attempt at the P vs. NP problem, and I regularly update my Stanford Encyclopedia of Philosophy entry on provability logic.

DG The topic of foundations of mathematics gives me a perfect bridge to my next question. Mathematical logic, at least in Frege's perspective, was born by distancing itself from psychology and 'psychologism'. In your work you seem to navigate the opposite route to reunite logic and the sciences involved with cognition. Can you tell our readers about your research path linking those two areas?

RV After my PhD I took up postdoc positions at Charles University of Prague and then at the University of Gothenburg, followed by a visiting assistant professorship at MIT. During those times I noticed two things that have been critical in shaping my future research. First, that at every conference on provability, interpretability and complexity I would go to, I would always meet the same colleagues. And second, that nobody outside that area would even read the papers we were writing about the foundations of arithmetic. The distance from real applications was quite large. So during my postdocs I was still writing up results about the foundations of mathematics, but I was preparing myself to move to another area, less removed from applications, and in which I would be able to interact with a larger variety of people. At that time a possibility presented itself to work in artificial intelligence (AI), at the Free University of Amsterdam, specifically on the applications of logic to AI. And that's when I got in contact with a much more heterogeneous environment than the one I was used to in the foundations of mathematics.

DG So we might say that it was the desire to engage with research questions bearing more closely on 'real world' applications that brought you into AI, and that has later become a gateway to cognitive sciences?

RV Yes, I was really trying to do something more useful for the world, so to say, and AI seemed a natural step in that direction. Then, when I came to the University of Groningen to work in the group on Cognitive Science and Engineering, many of my colleagues were cognitive scientists. So I became fascinated by their questions concerning human reasoning and cognition. Not surprisingly this interest was not immediately reciprocated. Some psychologists outside of our group would not hesitate to air their scepticism to me, citing famous experiments such as Wason's selection task: "Logic? How is that useful? It's well known that people don't reason logically!". But I was not convinced by that argument: the social version of the task is not really the same as the abstract logical version. And a few years later I was thrilled by the beautiful work by Keith Stenning and Michiel van Lambalgen on Wason's selection task, who highlighted the distinction between reasoning 'to' and 'from' an interpretation, and pointed to widespread misunderstandings in the communication between psychologists and their subjects in experiments on reasoning tasks, and therefore on the interpretation of such experiments. Problems that at first sight may look logically isomorphic may not at all be! That work convinced me of the important role that logicians can play in understanding how people really reason, and encouraged me to pursue my own research agenda in that area. In particular I became interested in reasoning in interactive contexts, where people reason about one another's mental states, and about the beliefs they have about others' mental states, that is, in higher-order theory of mind (ToM). I felt that in that area especially, insights from epistemic logic would be very useful. There were loads of articles in psychology on how children develop ToM, and usually those articles, at that time, did not distinguish between first-order ToM (that is, beliefs about somebody else's beliefs) and second- or higher-order ToM (that is, beliefs about somebody else's belief about somebody else's beliefs, and so on). They focused on when children develop first-order ToM, that is around the age of four, and then stopped there. While it is equally interesting to understand what happens later, when and how they acquire higher-order ToM, starting to be able to understand notions such as common knowledge involved in social constructs like promises.

DV So in this case the role of logic is a clarifying one. It helps to make explicit the theoretical background psychologists

may use to structure their experiments leading, ideally, to better experiments. Is that right?

RV Yes, but logic is also extremely useful for building computational cognitive models of agents, of direct relevance for AI. So, for instance, logic can help you in delineating all possible strategies that a reasoner can use in an interactive context such as a game. Logic can inform computational cognitive models, and in turn such models can provide predictions of human behavior in reasoning tasks: what they will decide, what they will attend to, as well as their reaction times during such tasks. In that way you can really marry logic, cognitive models and experiments.

DG So let me go further down this line concerning the role of logic in contemporary AI. Kant famously wrote "concepts without intuitions are empty, intuitions without concepts are blind". But recent stunning developments in AI seem to run in the face of that motto: where "big data" and "black box" AI appear to relegate concepts, and their analysis through logic, to the side line. Is there still a future for logic in AI research, and if so, what is it?

RV I think we, as logicians, may need to be a bit more assertive there. I have noticed statements by famous figures in AI saying that logic is dead and, true, deep learning has delivered impressive results in recent years. But I believe logic will again assert itself in this context as a key to the development of what is sometimes called explainable AI. State-of-the-art machine learning systems are examples of, as you say, a "black box" form of AI: they provide you with a decision but not with the reasons for you, as human user, to trust that decision at all. This is problematic, especially in view of the biases that such systems may inherit from hidden biases, cultural or even racial, that are present in their training data.

DG This is a promising view of logic as a means towards explainable AI and, if I understand you correctly, even as a way for data-driven AIs to overcome their potential biases through forms of introspective reasoning.

RV Yes, but I also believe that in order to pursue that vision we, as logicians – ha, now I do catch myself repeatedly saying "we, as logicians" – we should be more open-minded towards probabilistic and statistical forms of reasoning. Again, interdisciplinarity is inevitable!

DG This is much food for thought! And on the note of interdisciplinarity I want to conclude this interview with one last question. Interdisciplinarity is praised, in words at least, within several research communities and funding agencies, but we know that publishing a paper or obtaining funding with interdisciplinary ideas is extremely challenging. One is bound to upset reviewers from one disciplinary side or the other. What is your personal approach to pursuing successful interdisciplinary research, and do you have a word of advice for young researchers that would like to go down the path of interdisciplinarity?

RV What follows is of course only my personal view. But I think what is important is to keep a strong background in your own basic field and then to build a credible team around you. Take care that you are really an expert in one of the fields of relevance for your proposed research, and for the other fields, try to get strong experts on board. That's not something you can do if you only operate within the boundaries of one discipline. Again open-mindedness is essential to build such interactions. And definitely don't try to do everything yourself. It is simply impossible to be an expert on every single relevant area in complex interdisciplinary research.

DG Thank you very much Rineke!

News

Theoretical Aspects of Rationality and Knowledge, 24-26 July

The Sixteenth conference on Theoretical Aspects of Rationality and Knowledge (TARK 2017) took place from 24 July until 26 July at the University of Liverpool. This year, TARK was co-located with the workshop on Strategic Reasoning (SR 2017), which was held on 26 July and 27 July.

TARK featured invited presentations by Mike Wooldridge (Oxford), Edith Elkind (Oxford), Christian List (LSE), Pierpaolo Battigalli (Bocconi), Hans van Ditmarsch (LORIA) and Barteld Kooi (Groningen). Additionally, 37 contributed papers were presented. Of these contributed papers, 19 were introduced by a brief oral presentation followed by a poster presentation; the remaining 18 papers were introduced in a longer oral presentation.

In this space we cannot discuss each of these presentations individually, but we can briefly discuss some of the themes that were shared among many different presentations.

While not all presentations followed this pattern, many of them could be described as one of the following four types.

1. Presentations that *introduced* a new formal system, or proposed a *modification* to an existing system.

   Because TARK is about the *theoretical* aspects of rationality and knowledge, members of its community tend to like formal representations of action, knowledge, norms, et caetaera.

   There is no single system that can accurately represent every scenario related to rationality and knowledge. This means that, every so often, a new phenomenon is discovered that is not adequately representable in existing formalisms, prompting the definition of a new or enhanced system that can represent the phenomenon.

   New formalisms that were discussed at TARK 2017 include, but are not limited to, logics for knowing how/whether/what/why (Naumov and Tao, Wang), aggregation systems for group recommendations (Lev and Tennenholtz), and Boolean games with non-exclusive control (Belardinelli et al.).

2. Presentations about *technical results* related to existing formal systems.

   Once a formal system is defined, it becomes important to study its technical properties, in order to gain a better understanding of both the formal system and the phenomenon that it is intended to model.

   A non-exhaustive sample of the technical results presented at TARK 2017 includes a number of impossibility results in social choice theory (Peters, Lev and Tennenholtz), a

number of (non-)axiomatizability results in quantified update logics (Galimullin and Alechina, Kuijer), decidability results for an extension of the 'logic of gossips' with common knowledge (Apt and Wojtczak), and yes/no-trade theorems (Hellman and Gizatulina).

3. Presentations that *translate* between two or more existing formal systems, showing the ways in which the systems are related.

   TARK brings together scholars from different fields, including (but not limited to) artificial intelligence, computer science, economics, logic, linguistics, philosophy and psychology. It is therefore not surprising that the TARK community uses a wide variety of different formal frameworks.

   As a result, comparisons between different systems are of particular interest. Some examples of these kind of comparison results presented at TARK 2017 include comparisons between Type Spaces and Probability Frames (Bjorndahl and Halpern), between Bayesian reasoning and narration-based reasoning in law (Urbaniak), and between Choice Theory and Deontic Logic (Dietrich and List).

4. Presentations that *apply* a known formal system to gain insight into a particular phenomenon.

   Despite the theoretical focus from TARK, some work also show how known formal tools offer new insight when used to model a specific phenomenon. Two examples of such new applications proposed at TARK 2017 are a epistemic-temporal analysis of the blockchain protocol (Halpern and Pass), an epistemic approach of psychological games (Jagau and Perea), and a judgment aggregation analysis of liquid democracy (Christoff and Grossi).

Zoé Christoff and Louwe Kuijer
Amsterdam and Liverpool

Formal Models of Scientific Inquiry, 18-19 July

The conference on Formal Models of Scientific Inquiry took place on 18-19 July at the Ruhr-University Bochum. It was organized by AnneMarie Borg, Dunja Šešelja, Christian Straßer – members of the Research Group for Non-Monotonic Logic and Formal Argumentation (NMLFA) at the Ruhr-University Bochum, and Vlasta Sikimić from the University of Belgrade. The keynote speakers were Jason McKenzie Alexander (London School of Economics and Political Science), Gregor Betz (Karlsruhe Institute of Technology) and Leah Henderson (University of Groningen).

Philosophical discussions about social aspects of scientific inquiry are increasingly turning towards the utilization of formal models. Researchers from different continents met and discussed different formal models used for the optimization of scientific reasoning. The conference targeted both state-of-the-art research in formal models of scientific inquiry such as agent-based models and data-calibrated models, as well as methodological discussions about designing informative models. Participants addressed the question of when models of scientific inquiry are informative for science policy. One of the most important insights was that empirical calibrations are needed for a better interpretation and application of formal models. The intention is that the conference becomes a biannual event.

Gregor Betz (Karlsruhe Institute of Technology) opened the conference with his talk on probabilistic confirmation theory with an argumentation-theoretic flavor. This setting was applied for establishing the unity of practical and theoretical reasons. Contributed talks included the citation based analysis of climate science, formal models of discrimination in science, the game theoretic approach in modelling social mechanisms in science, the examination of the variety of evidence thesis, formal argumentation frameworks, etc.

Members of the NMLFA research group, consisting of AnneMarie Borg, Daniel Frey, Dunja Šešelja and Christian Straßer, presented their agent-based model of scientific inquiry which employs an abstract argumentation framework. Their results suggest that a high degree of information flow among agents is epistemically beneficial. Remco Heesen (University of Cambridge) presented his joint work with Jan-Willem Romeijn (University of Groningen) about biases in editorial policies of scientific journals, ending his talk with a bold suggestion that we might be better off without a peer review system.

Slobodan Perović (University of Belgrade) presented the results of his group which included big data analysis of project efficiency in the Fermi National Laboratory, and the analysis of the reliability of citation metrics in high energy physics. Due to the inductive nature of the research in high energy physics, the citation metric is a reliable measure of project efficiency in the field. Michael Thicke (Bard College) presented his work on citation analysis in climate science. He emphasized the importance of determining how scientists actually interact. Though informative for modelling, data about scientific interaction are challenging to obtain and interpret. Thicke also addressed these difficulties in his talk.

On the second day, in her invited talk, Leah Henderson (University of Groningen) showed how Hierarchical Bayesian models can be used for understanding the interplay between empirical support of specific models and the general structure of a scientific theory. More talks on Bayesian modelling followed.

In the final conference talk, Jason McKenzie Alexander (London School of Economics and Political Science) discussed criteria for an informative model of scientific inquiry, pointing out that some empirical calibration is necessary. Though models come in a degree of idealization, they need to meet certain criteria to be informative. One important criterion is the empirical calibration of a model. These considerations concluded the event.

The conference was supported by the Alexander von Humboldt Foundation and the Ruhr University Research School Plus.

Vlasta Sikimic
Philosophy, University of Belgrad

What's Hot in ...

(Formal) Argumentation Theory

A question posed on my recent (academic) travels, and that is regularly raised in workshops and conferences, is what the relationship is between argumentative characterisations of non-monotonic inference, and the panoply of formal models of argumentation that study various kinds of relations between arguments. As described in the June'17 edition of the Reasoner, the former considers arguments constructed from a set of logical formulae Δ. These arguments are related by binary attacks, where (for example) an attacking argument concludes some α that negates a premise or (possibly intermediate) conclusion in the attacked argument. The sets of justified arguments in the framework consisting of arguments and attacks, are then evaluated. The claims of justified arguments are then shown to correspond (for a number of established non-monotonic logics) to the inferences defined directly over Δ.

Since Dung's seminal work initiating research into these logical instantiations of argument frameworks, substantial research has been invested into the evaluation of arguments in frameworks relating arguments in ways other than by binary attacks. For example, numerous works study various kinds of support relations between arguments, and 'collective attacks' by sets of arguments on individual arguments. Most of these works do not study the generation of these more 'exotic' frameworks by instantiating sets of logical formulae. Rather, they usually cite natural language examples of human uses of argument by way of motivating these various species of relations amongst arguments. This then raises the question of how one should validate the 'correctness' of the outcomes of evaluation mechanisms specified for these exotic frameworks, and with respect to what is such correctness to be judged ? One answer might be that since these frameworks provide descriptive accounts of human argumentation, empirical studies of human evaluation of arguments should provide standards by which one judges the correctness of the proposed evaluation mechanisms. However, this does not suffice if formal models of argumentation are to provide normative guidance for human reasoning and debate. With this in mind, one suggestion would be to recognise how additional relations warranted by human uses of argument arise because of the ways in which human argumentation deviates from formal logic-based argumentation. Formal models can then be deployed to provide prescriptive guidance.

For example, humans often make use of incomplete arguments (so called *enthymemes*). Typically, an argument claiming that Tweety flies because Tweety is a bird, would not be countered by a complete argument of the form "But Tweety is a penguin and penguins do not 'fly". Rather, a more typical reply would be the enthymeme "But Tweety is a penguin"; a Searle type indirect speech act in which the use of the word "But" indicates that the enthymeme is an attack, and that there is an implicit shared understanding that the missing information is the rule "penguins do not fly". This analysis accounts for the use of collective attacks in human argumentation, in which multiple arguments $A_1, \ldots, A_n$ (respectively claiming $\alpha_1, \ldots, \alpha_n$) collectively attack an argument B. Motivating examples typically suggest that in fact, there is an implicit missing rule with antecedents $\alpha_1, \ldots, \alpha_n$ and a consequent β that negates some premise or (intermediate) conclusion in B (e.g., B = "the information was obtained by unethical means and so we should not publish" collectively attacked by A_1 = "the information is available on the internet" and A_2 = "the information is in the public interest"). Moreover, there may be some ambiguity as to what is the identity of the β in question (i.e. precisely which premise or conclusion in B is being targeted). Prescriptive guidance would encourage the human interlocutors to clarify the identity of β ("do you therefore mean that we should publish or that we can overlook the unethical provenance of the information ?"), and thus the missing rule. Furthermore, the arguments $A_1, \ldots, A_n$ can then be extended with the identified rule to define a single argument A that claims β and that binary attacks B. The rational outcome of the exchange is then evaluated by reference to a framework consisting of arguments related only by binary attacks.

Sanjay Modgil
Informatics, King's College London

Medieval Reasoning

In any research area there are topics that, at some point in time, are the talk of their day. Sometimes their popularity is due to their novelty and to reasons that are internal to the field: maybe a breakthrough has recently been made; perhaps it is an entirely new subject yet to be explored, and so on. Sometimes it has more to do with extrinsic factors – be they social, institutional or merely accidental. In most instances, however, there are both internal and external reasons for a topic to get hot and start popping up everywhere in journals, calls for papers and applications, among colleagues talking shop, and at every other conference. Conferences – and in particular the one "big" conference that every field has – play a decisive role in anointing "The Topic" that everyone is going to focus on for a while. For those dealing in the history of medieval logic and related subjects, the "hot conference" is undoubtedly the biennial European Symposium on Medieval Logic and Semantics. While "symposium" usually refers to a very specialised conference in some particular subfield – as the ESMLS certainly is – in this instance the title is literally well deserved: in 1973, as Sten Ebbesen recalls, a small group of scholars (including Miecislaus Markowski, Jan Pinborg, and Lambert Marie de Rijk) decided to meet in a café in Warsaw for "a detailed study of Polish vodka" (sic!) and of medieval logic. A couple of years later the Symposium reconvened in Copenhagen, this time in a more formal fashion; the resulting volume ("The Logic of John Buridan" ed. J. Pinborg, 1976) was among the first major contributions bringing Buridan into the spotlight as one of the most interesting medieval logicians – Hubien's Latin edition of Buridan's "Treatise on Consequences" was published

that same year. Since then, the ESMLS has become a regular appointment for the old guard and the next generation of young scholars to mingle and discuss the hot topic of the day in medieval logic and semantics, occasionally with a more pronounced metaphysical twist – for example, the last couple of editions have been, respectively, on Theories of Relations (Cambridge, 2014) and Mereology (Pisa, 2018). Next June in Düsseldorf, the focus is going to be on "The Continental and British Traditions of Medieval Logic Revisited". This split between a Continental and a British tradition is, in a sense, the medieval version of the Analytical/Continental divide, at least insofar as it is a classificatory approximation resting upon some deep differences in approaches and philosophical interests. As is often the case with such classifications, these differences are neither uniform nor constant: the philosophical discussions between Paris and Oxford in the 12th century are not the same as in the 14th; and most importantly, these neatly labelled boxes are often too narrow to really account for the influences and hybridisations which, at a closer look, are evident on either side of the fence. The Continental vs. British divide and its limits might seem to be somewhat old news: in 1979 the topic of the ESMLS was "English Logic and Semantics from the 12th Century to the Time of Ockham and Burleigh" (Leiden); in 1980 it was "English Logic in Italy in the 14th and 15th Centuries" (Rome); and in 1983 it was "The Rise of British Logic" (Oxford). It is old news indeed. However, on the one hand, as it is too often the case with medieval "old news", there are still many pieces missing (unedited texts, unidentified authors, unintelligible or misunderstood theories...) for the picture to be complete. On the other hand, the pieces we do have do not always fit well together: our old news is probably at least partially fake and in dire need of a thorough reassessment. And this is quite a hot topic, if we want to understand better not only what was going on in medieval logic, but also how these discussions and intellectual practices developed. Overall, it looks like the 2018 ESMLS is going to be a pretty hot conference, as usual. The CFP is now open! Get in touch with the organisers Christian Rode (christian.rode@online.de) and Christoph Kann (kann@phil-fak.uni-duesseldorf.de), if you would like to attend, contribute – or take part in a detailed study of German ale.

Graziana Ciola
Philosophy, Scuola Normale Superiore, Pisa

Uncertain Reasoning

I was at a very interesting workshop in Bochum recently on "Experience and Updating", which prompted me to think about evidence and "evidentialism". Call evidentialism the view that your belief state is rational if and only if it is justified by your total evidence. So what is it for evidence to justify a belief state? Well, what *is* evidence?

In probabilistic contexts, one often hears people say things like "imagine that E is evidence for H" where E and H are taken to be propositions: elements of the algebra of the probability space. And what is it for E to be evidence for H in this context? Presumably (other things being equal) for $p(H|E)$ to be greater than $p(H)$. (If you have some other more sophisticated measure of evidential strength, plug it in here: nothing I say will depend on that...) So E's being evidence for H is simply a property of the prior probability function, namely that $p(HE) > p(H)p(E)$.

But in most contexts we'd want to say that certain probability functions "get it wrong" about the evidential relationships. Consider trying to guess how likely it is that a coin of unknown bias will land heads on the 11th toss. And consider doing so on the basis of the evidence that the first 10 tosses contained 8 heads. There's nothing probabilistically wrong with having the probability for heads on toss 11 to go *down* on learning this evidence unless we make further assumptions. Obviously, these assumptions seem natural to us (for example, it's natural that the coin toss events should be exchangeable), but where do these judgements of naturalness come from? Certainly not from the theory of probability on its own. So "what is evidence for what" should be something that is given exogenously. In short, I think that we need a *theory* of evidential support that is somehow prior to the assigning of priors to propositions.

Here's another reason to think that such a theory is needed. In probabilistic contexts, we almost always suppose that our total evidence is consistent, and that new additional evidence is consistent with already gathered evidence. But our acquisition of evidence is fallible, and it seems like there is room for a reasonable theory of how to accommodate new data that contradicts our old evidence. Perhaps some form of AGM revision (Gärdenfors, Knowledge in Flux, 1988 MIT Press), or some kind of Pollock-style defeasible reasoning (Pollock, Cognitive Carpentry, 1995 MIT Press)? Again, I think we need a theory of evidence that is somewhat independent from "bayesianism".

There has recently been a lively debate about whether probabilistic conditioning (or something like it) can accommodate "defeating" evidence (see, for example, Weisberg "Commutativity or Holism? A Dilemma for conditionalisers", 2009 BJPS). I feel like this debate is premature. Let's work out what evidence is, how it constrains or justifies rational credences (and rational changes in credence) and only then can we see whether conditioning is an adequate theory of change in belief. I have the same feeling about the recent discussions of whether "The Principal Principle Implies the Principle of Indifference" (Hawthorne et al, 2015 BJPS). I feel like we need a better handle on a theory of evidence to be able to make the claims that Hawthorne and co make about whether certain propositions are admissible or not.

It seems that the time is right to step away from these particular debates and think more generally and more carefully about what evidence is, how evidence changes over time, and how that constrains rational belief.

Seamus Bradley
Philosophy, University of Tilburg

Evidence-Based Medicine

A jailed breast surgeon last month had his sentence increased. He was previously convicted for performing unnecessary or inadequate surgeries, including the so-called cleavage-sparing mastectomy. The idea behind the cleavage-sparing mastectomy is to minimize any change in the appearance of a woman's cleavage following surgery for breast cancer by leaving behind

some breast tissue. The procedure is unorthodox and unregulated: It is not recommended by any health care guidelines because it is believed to increase the risk of breast cancer recurrence compared to standard mastectomy. Interestingly, the history behind those guidelines looks to provide an example of the importance of combining evidence from clinical trials with evidence of mechanisms.

The standard surgical treatment for breast cancer for much of the twentieth-century was the radical mastectomy popularized by Sir William Stewart Halsted. Halsted was widely acclaimed as a skilled surgeon. This is particularly impressive given that he was typically under the influence of narcotics. He had developed an addiction to cocaine after experimenting on himself in order to establish its local anaesthetic properties, and this addiction was cured only by getting addicted to morphine instead. However, it was his early adoption of anaesthetics in breast cancer patients which allowed him to develop the radical mastectomy by permitting more extensive and precise surgical procedures. Radical mastectomy was based on the *anatomic* theory of the mechanisms of breast cancer, namely, that cancer cells from the breast initially spread only as far as local lymph nodes where they are trapped for some time before being spread throughout the body. It was therefore thought that a local but large operation was necessary to remove the cancer before it had spread throughout the body: Radical mastectomy involved removing the breast, as well as nearby lymph nodes and chest muscles.

In the later half of twentieth-century, the anatomic theory of the mechanisms of breast cancer was challenged by Bernard Fisher. As a newcomer to the field, Fisher had made a couple of observations about the state of breast cancer research. On the one hand, he was impressed by 'how little information there was related to the biology of breast cancer and what a lack of interest there was in understanding the disease'. As a result, Fisher himself conducted a number of laboratory investigations into the mechanisms of breast cancer. On the basis of these investigations, he proposed an alternative *systemic* theory of the mechanisms of breast cancer, and he pointed out the implications of such a theory for the surgical treatment of breast cancer:

> From those laboratory studies, we formulated an alternative hypothesis—that breast cancer was a systemic disease in that tumor cells were likely to have been disseminated throughout the body by the time of diagnosis and that more expansive locoregional therapy was unlikely to improve survival. As a consequence, less radical surgery was likely to result in similar outcomes to those obtained following radical mastectomy.

On the other hand, Fisher had also acknowledged the need for randomized clinical trials in testing proposed surgical interventions for breast cancer. Accordingly, he carried out a clinical trial in order to complement his laboratory investigations: 'In that study, patients were randomly assigned to receive a Halsted radical mastectomy, a total (simple) mastectomy, or a total mastectomy followed by radiation therapy. In the 25-year follow-up of that study which was published in 2002 in the *New England Journal of Medicine*, the results from [the study] continued to indicate that there was no difference in overall survival, disease-free survival, or survival among the three groups of women'. Arguably, it was only this combination of evidence of mechanisms together with evidence from clinical trials which led to guidelines shortly after recommending against radical mastectomy.

Michael Wilde
Philosophy, Kent

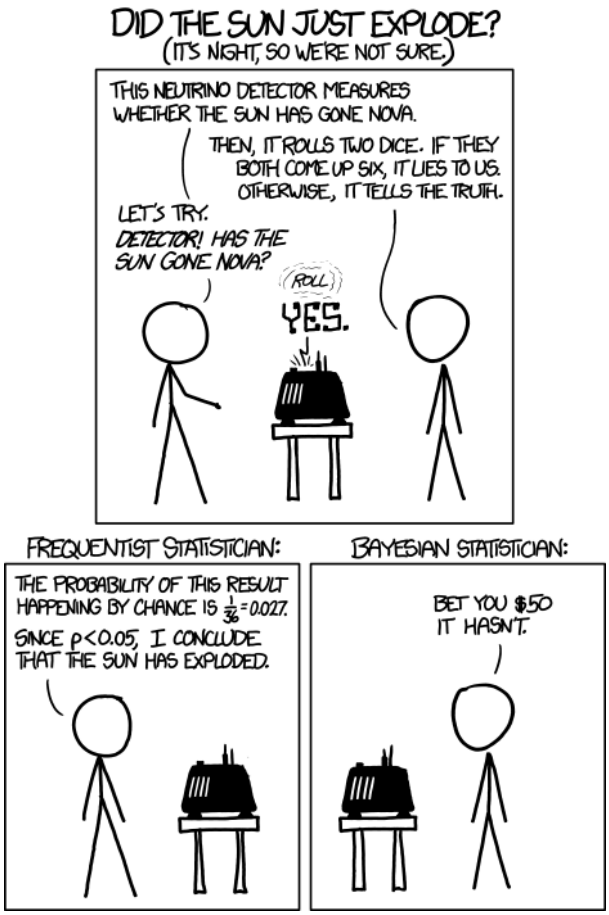

Events

September

TNF: Teleosemantics and the Nature of Functions, Bielefeld University, 7–8 September.

SubBay: The Fifth Subjective Bayesian Meeting, The Open University, Milton Keynes, 12 September.

EpiVaC: Epistemic Vice and Corruption University of Nottingham, 12–13 September.

MPR: Modeling Physical Reality, Salzburg University, 13 September.

CauFW DA: Causality, Free Will, Divine Action, Vienna, 13–15 September.

AiC: Attitudes in Context, University of Regensburg, 14–15 September.

EPINON: Epistemology in Ontologies, Bozen-Bolzano, Italy, 21–23 September.

EpiRat: Epistemic Rationality: Conceptions and Challenges, Barcelona, 21–23 September.

BCC: Bridges Causality Conference, University of Warwick, 28–29 September.

AG&CO: Workshop: Agency and Control, Copenhagen, 29 September.
RRAIA: Workshop on Reasons, Rationality, and Intentional Agency, London School of Economics, 29–30 September.

OCTOBER

HISTORY AND PHILOSOPHY OF COMPUTING: Brno, 4–7 October.
NTIE: New Trends in Epistemology, Hamburg, 5–7 October.
MARU: Moral and Rational Uncertainty, University of Reading, 9 October.
RLHRC: Representation Learning for Human and Robot Cognition, Bielefeld University, Germany, 17 October.

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

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.

JOBS AND STUDENTSHIPS

Jobs

POST DOC: in Probabilistic Machine Learning, National University of Singapore, deadline open.
POST DOC: in History and Philosophy of Science, University of Edinburgh, deadline 21 September.
PROFESSORSHIP: in Theoretical Philosophy, Konstanz, Germany, deadline 28 September.
PROFESSORSHIP: in Statistics, University of Bath, deadline 16 October.
READER: in Statistics, University of Bath, deadline 16 October.

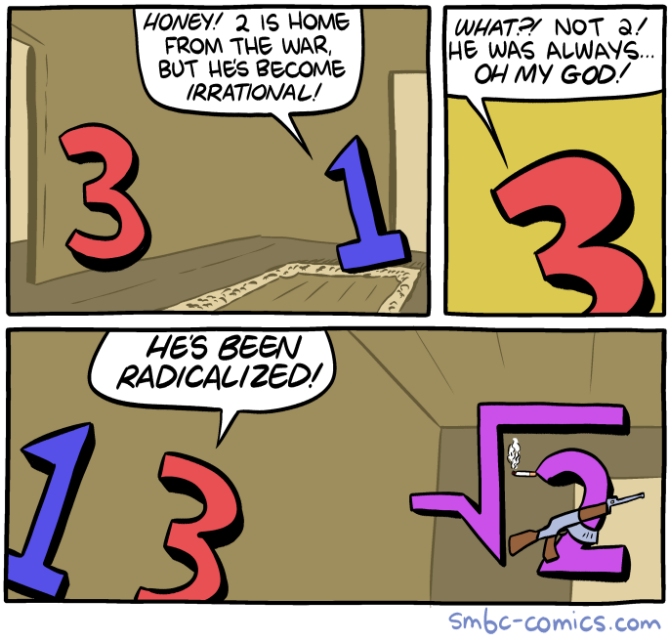

LECTURER: in Applied Mathematics, University of Bath, deadline 16 October.
ASSISTANT PROFESSOR: in Metaphysics and Epistemology, University of Toronto, deadline 1 November.

Studentships

5 PHD'S: in Philosophy, University of Milan, Italy, deadline 4 September.
PHD: in Epistemology/Philosophy of Mind, University of Fribourg, Switzerland, deadline 30 September.