The Reasoner

Volume 13, Number 2
February 2019

thereasoner.org
ISSN 1757-0522

Contents

Guest Editorial

Dear fellow Reasoners,

I'm one of those people who think that what is rational to believe is subject to evidential constraints which govern our epistemic lives. For example, I have very good evidence that, if I bet on the number 13 in Roulette, then I will most probably lose the chips I put down. I should hence definitely not assign more than 80% probability to the proposition/event that the number 13 will be the lucky number the next time the wheel is spun. I invested a lot of time in trying to answer this question, which probability exactly I should assign to events on epistemological grounds. Today, I believe that their is no single correct answer, although Maximum Entropy Methods for answering this question almost always provide one good answer.

It is hence a great pleasure that Professor Gabriele Kern-Isberner from the TU Dortmund, a fellow Maximum Entropy aficionado, shares her thoughts on uncertain reasoning in general and Maximum Entropy in particular. Not only did I find her views on reasoning intellectually stimulating, as a father of three, I also take her experience with teaching programming to children as good parenting advise. This then leads me to the interview I gave to *The Reasoner* (2018, Interview with Jürgen Landes, The Reasoner, 93-96) last November, in which Hykel Hosni asked me how to get the public interested in reasoning. Somewhat naively, I suggested to engage children on grounds of their natural inquisitivity. I look forward to follow my own suggestion by teaching programming.

In this interview, I also said that "we are just not born with innate interest in reasoning. What a shame!". Gaby's observations do not quite conform with my claim. Looking back, I still think that we are not born with an innate interest in reasoning proper, but – at the very least – I ought to have added that this claim solely rests on my own observations. On a more positive note, Gaby points to transferring reasoning methods to first order logic as an important challenge. Incidentally, I happen to work on a Maximum Entropy application to first order logic.

Before I treat you to this month's interview, I would like to state my appreciation of *The Reasoner* and the intellectual stimulation it has provided me over the years and – much more recently – parenting advise. As a reasoner and contributor to *The Reasoner* (2018, What's Hot in Mathematical Philosophy: Pirate Games, 41-42 and 2018, L&P-updating - All Bets Are Off, 10), I had a sense of excitement when the new editor Hykel

Hosni(2017, Editorial, The Reasoner, 22-23) announced the introduction of novel ways reasoners could contribute. Unfortunately, the number of such contributions has not reached my (too lofty?) expectations. Let me thus add my name to the list of people Hosni(2017, Editorial, The Reasoner, 91) who would like to read more novel features.

JÜRGEN LANDES
Munich Center for Mathematical Philosophy

FEATURES

Interview with Gabriele Kern-Isberner

JL: Could you please tell us how you got interested in reasoning?

GKI: Actually, I have always been very interested in knowledge and reasoning, already when I was a child. I've always been fascinated by knowledge and considered it as my capital and best warrant for a good and satisfying life. Not that I liked school, but I was eager to learn, preferably in an autodidactic way, and using my knowledge (which is reasoning) for whatever goal. Much later, my first contact to artificial intelligence and knowledge representation was through my children: I found the book *Mindstorms* by Seymour Papert in a library and couldn't stop reading his ideas on how computers can enhance learning and reasoning capabilities in children. Children are really experts in learning, especially when they are very young. I liked watching my children playing, or exploring something, one can very easily recognize various ways of basic nonmonotonic reasoning, belief revision, and learning from small samples in children's activities, and they combine all this in an effortless way. They count and compare, guess, try, and revise, and start once again, always improving their world view and their competencies. This is my model of reasoning, of intelligence, a seamless integration of (qualitative and quantitative) learning, reasoning, and revising, guided by few basic principles which show different facets in different contexts. Probability theory is a good example for that – using conditional probabilities, beliefs can change drastically. So, probabilities are highly nonmonotonic but obey few axiomatic principles. Beyond that, I think, the principle of maximum entropy is also one of these basic principles for reasoning tasks (as I will explain later on).

JL: What are the important problems in your area of research?

GKI: As my main areas of research, I would consider uncertain and nonmonotonic reasoning, and belief revision. Qualitative nonmonotonic reasoning and belief revision are usually very formal, in parts even philosphical, they are lacking successful applications, algorithms, and implementations, so this would be an important problem. Probability theory has experienced an enormous success story since the 80s of the last century but choked many other interesting approaches to uncertain reasoning. If one looks into the proceedings of the UAI conferences (Uncertainty in Artificial Intelligence) in the 80s and 90s, they show an impressive breadth of approaches with innovative ideas, nowadays it is mainly about statistical work. As usual, this problem has two sides: On the one side, in my opinion, it would be a good idea if qualitative reasoners look into techniques and algorithms more closely that have been developed for probabilistic reasoning, on the other side, probabilistic and statistical reasoners should be concerned more by the qualitative, formal, cognitive dimensions of their inference methods. Traditions in qualitative vs. quantitative reasoning are very different, there are lots of controversies. Although even very prominent researchers like Pearl and Halpern published beautiful works connecting qualitative and quantitative aspects of reasoning, there is still a remarkable gap, in particular, if one considers belief revision and machine learning. Although both areas deal with adopting and integrating knowledge, differences between the communities are huge, the communities are nearly disjoint. The principle of minimum cross-entropy is a powerful belief change method for probabilities and used in various forms for machine learning but hardly known and understood in the belief revision community. So, on a meta level, a big challenge would be to overcome fragmentation within the field of reasoning, where we have divisions into different subcommunities, even within the subareas. In February, I will co-organize a Dagstuhl seminar where we aim at bringing subcommunities of nonmonotonic reasoning and belief revision together again, trying to re-gain a more global picture and finding innovative and relevant perspectives for the future. There are also initiatives from the machine learning side to connect more closely to logics (there is another Dagstuhl seminar on this topic in 2019, organized by Kristian Kersting and others), so currently, there seems to be a good time to re-try to close the gap.

On a more methodological level, a huge challenge is to generalize and transfer reasoning methods which have been developed for a propositional base logic, to the first-order case. Actually, this is much more than just generalizing and transferring – going for first-order, or relational approaches in uncertain reasoning opens up new dimensions of reasoning because first-order interpretations are much more complex than propositional ones, even if one uses Herbrand semantics. Here, subjective (degrees of belief) and objective (statistical view) of probabilities clash, and one has to reconcile both views. The huge amount of approaches to first-order probabilities that have been put forward within the last two decades illustrates the variety of the topic. For probabilities, it is a very active field of research, but for qualitative reasoning methods, first-order approaches are rare to date, except for the big and successful field of extended logic programming which makes use of first-order representations and basically relies on a Herbrand semantics. Applications of uncertain reasoning to description logics which are basically classical-logical are a major challenge currently and could result in significant progress towards the aim of bringing qualitative and quantitative reasoning together.

JL: What are you currently working on?

GKI: For the last two decades, I have been working on transferring basic ideas from the entropy principles to qualitative reasoning, and this is still a focus of my ongoing research. In the late 90s of the past century, I published an axiomatic characterization of the principle of minimum cross-entropy, and one of those axioms was the so-called principle of conditional preservation (2004, "A Thorough Axiomatization of a Principle of Conditional Preservation in Belief Revi-

sion", *Annals of Mathematics and Artificial Intelligence*, 127-164) This is an invariance property that minimum cross-entropy and maximum entropy distributions satisfy, and which consists of a purely algebraic part and a semantic (i.e., probabilistic) part. By changing the semantic part, this principle can be transferred to other semantic frameworks like possibility theory, or Spohn's ranking theory which can be considered as a logarithmic abstraction of infinitesimal probability theory. We have been working a lot with these ranking functions because many people find them quite intuitive, even students like working with them, and they are quite a perfect semantic environment to implement some basic properties of the entropy principles for qualitative reasoning, in particular, this principle of conditional preservation. On an informal level, this principle ensures that conditional dependencies in the prior distribution, or ranking function, are preserved as much as possible when new conditional dependencies are learnt. Therefore, it helps solving a very advanced belief change problem: How should prior knowledge be changed when the new information consists of a set of conditional beliefs, to be adopted or to be given up. Quite recently, we have explored this principle quite thoroughly for different qualitative change operations like revision and contraction (as a form of forgetting) and were able to show that it yields many other approaches as specializations. At KR 2018, I presented a purely qualitative principle of conditional preservation from which axioms of a seminal paper by Darwiche and Pearl on iterated belief revision can be derived. But my feeling is that this qualitative principle of conditional preservation has many more interesting applications in qualitative reasoning and maybe can solve some open problems, hopefully give rise to novel formal general properties of nonmonotonic reasoning.

For probabilistic reasoning on maximum entropy, we have transferred the principle of maximum entropy to first-order probabilistic knowledge bases, and we are currently exploiting the algebraic grounds of the maximum entropy principle (relying on this principle of conditional preservation) for developing algorithms which can solve the lifted inference problem in first-order probabilistic reasoning efficiently.

JL: 40 years ago, E. T. Jaynes wrote his paper: Where do we stand on Maximum Entropy (in The Maximum Entropy Formalism, MIT Press, 1978, Chapter 1, 15-118). what are the main developments since then? Where do we stand on MaxEnt, today?

GKI: This is definitely a very broad question, and I can only say something on MaxEnt in computer science and parts of mathematics, and even here, I think I don't have a good overview. From my perception, MaxEnt (and related principles) has nowadays similar problems as those that I described above for reasoning in general: the field has much diversified and is fragmented. Within the past few years I attended two workshops on entropy that Laura Martignon organized in Ludwigsburg, and it was very informative to see on which aspects of entropy people are working, however, it was nearly impossible to find a common base on which we could discuss entropy. Many different variations of entropy or similar measures were discussed and compared, it was all very interesting to see these ideas, but I could not draw any satisfactory reconciliation from this. I would like to see a kind of catalogue of formal (preferably basic) properties for probabilistic entropy-like reasoning methods according to which probabilistic reasoning methods can be classified. Based on such properties, one can even define more clearly what should be called an entropy, and what not. Such formal properties have been presented by Shore and Johnson, and most prominently, by Jeff Paris, but they do not seem to be in the focus of probabilistic reasoning currently. What I liked a lot was Jeff Paris' paper on connecting MaxEnt reasoning to commonsense reasoning (1998, "Common Sense and Maximum Entropy", *Synthese*, 75-93). Beyond all these statistical and numerical aspects of entropy, this commonsense, or cognitive perspective of MaxEnt should be explored better.

JL: What are your views on the impact the reasoning community/ies have?

GKI: Indeed, this is a crucial question with quite a sad answer: The impact of the reasoning community is low, too low, in particular, at the time being with all this AI hype, in particular, when compared to the success of deep learning currently. The general view on reasoning, even in computer science, is mostly restricted to either classical-logical, or to probabilistic-statistical frameworks, and then reasoning has to be efficient and economically successful. Very often, reasoning is not understood as a methodology in its own, but just as an auxiliary instrument which has to do its work. Every one is talking about knowledge and information, but what are both without reasoning methods? At last KR 2018, we had a plenary discussion on the role of reasoning, and its outcome revealed nearly the same depressing picture. If anyone has a good idea of how to change this, this would be very welcome!

JL: What role does MaxEnt play for reasoning in general?

GKI: I think that the entropy principles (in their original form in the work of Jaynes) play a most basic and versatile role for reasoning, not because of their statistical properties but because of their conditional-logical properties. I mentioned the principle of conditional preservation as a main building block of reasoning at optimum entropy which means that the entropy principles process conditional dependencies in a most adequate way. While conditional independencies are crucial for Bayesian networks, conditional dependencies are the lines which reasoning at optimum entropy follows. This observation has two important implications: First, reasoning at optimum entropy has a conditional-logical quality that is rarely found in other reasoning methods, and that ensures formal properties of reasoning and belief change on highest levels. Second, conditionals (as three-valued entities in the sense of de Finetti) are basic building blocks of our (commonsense and specialist) knowledge, even small children are able to understand, utter, and process them, as in "You promised to give me gummy bears when I'm nice, and I was nice!" - if nice then gummy bear, so having been nice, the child expects to get gummy bears. Therefore, the entropy principles (and their qualitative counterparts based on ranking functions) connect good logical quality with the ability to represent and process commonsense knowledge, this should make them best candidates (at least from a theoretical point of view, implementations can be nasty) for any reasoning task in AI. However, these two crucial aspects of entropy principles, the formal-logical one and the common sense one, seem to be less well-known. Moreover, entropy principles help to close gaps: As described above, their basic axiomatic properties can be transferred to other semantic frameworks, and they can realize both nonmonotonic reasoning and belief change in one methodological framework (as the principle of maximum entropy can be considered as a special case of the principle of minimum cross-entropy). That's why I think that the entropy principles implement basic principles of reasoning in a seamless way, and are really valuable objects of research.

News

Calls for Papers

Knowing the Unknown: Philosophical Perspectives on Ignorance: special issue of *Synthese*, deadline 20 February.
Hybrid Data and Knowledge Driven Decision Making under Uncertainty: special issue of *Information Sciencesl*, deadline 30 February.
Computational Modeling in Philosophy: special issue of *Synthese*, deadline 1 March.
Thought Experiments in the History of Philosophy of Science: special issue of *HOPOS*, deadline 31 March.
Folk Psychology: Pluralistic Approaches: special issue of *Synthese*, deadline 15 May.
Imprecise Probabilities, Logic and Rationality: special issue of *International Journal of Approximate Reasoning*, deadline 1 June.

Dissemination Corner

The Logic of Conceivability

Recent output Since our last post in *the Reasoner*'s Dissemination Corner, the Logic of Conceivability Project (LoC) conceived and realised a number of papers. Franz' 'Taming the Runabout Imagination Ticket' appeared in *Synthese* (2018), 'The Theory of Topic-Sensitive Intentional Modals' in *The Logica Yearbook 2018*, and his joint work with Peter, 'Knowability Relative to Information,' in *Mind* (2018). Aybuke's 'A Topological Approach to Full Belief,' in collaboration with A. Baltag, N. Bezhanishvili, and S. Smets, has appeared in Journal of Philosophical Logic (2018), while Karolina's joint work with Niels Skovgaard-Olsen, Peter Collins, Ulrike Hahn, and Karl Christoph Klauer is forthcoming in *Cognitive Psychology*. Finally, Franz' book with Mark Jago, *Impossible Worlds*, will appear in 2019 (Oxford University Press).

The LoC researchers have meanwhile turned their attention to the notion of relevance between a conditional antecedent and its consequent. Here I survey some of the developments that sparked our interest in this phenomenon.

The Relevance of Relevance One of the aims of LoC is to study how people reason when they imagine non-actual situations, that is, when they think about what might happen or what might have happened. Among others, this kind of reasoning plays an important role in our production and interpretation of indicative conditionals, such as: "If you publish in good journals, you will get tenure" or "If we do not reduce our green house gas emission, the climate change catastrophe is inevitable." One aspect of the interpretation of conditionals that became LoC's focus is the connection between a conditional's antecedent and its consequent. This connection can be understood in various ways, for instance, as an evidential or inferential relation, as a causal or explanatory link, or as probabilistic relevance. There remains, however, a more fundamental question pertaining to the nature of the connection: does it belong to the (broadly construed) semantics of a conditional or is it merely a pragmatic aspect of its meaning?

On the vast majority of theories of conditionals, the connection plays no role in determining the truth value or the acceptability value of a conditional. On those theories, if the significance of the intuition that the antecedent and consequent should be connected is acknowledged at all, it is considered to be a purely pragmatic phenomenon. Nonetheless, it is not an entirely new idea that the connection belongs to what is literally said: the conventional, semantic content of a conditional, and hence that it contributes to its truth or acceptability conditions. The view that the consequent should be *inferrible* from the antecedent has been advocated, for instance, already in *A System of Logic* by John Stuart Mill (1843). The 20th century has witnessed attempts to capture the connection between antecedents and consequents in a formal system, such as relevance logics or Barwise and Perry's situation semantics, but none of these became mainstream. What triggered a new wave of interest in the status of the connection between antecedents and consequents have been recent developments in cognitive science.

The first bits of evidence for the "inferential" approach to conditionals can be found in the work by Douven and Verbrugge (2010: 'The Adams Family', *Cognition*), who have drawn directly from the empirical linguistics, such as, for instance, the corpus based analysis of conditionals in English by Declerck and Reed (2010: *Conditionals: A comprehensive empirical analysis*, Mouton de Gruyter), where different types of conditionals sentences are characterised in terms of different kinds of relations connecting their antecedents and consequents. Taking the notion of an inferential conditional as their starting point, Douven and Verbrugge investigated how different types of the inferential link between antecedents and consequents affect people's acceptability and probability ratings. More specifically, they investigated different versions of the so-called Adams Thesis, according to which the *acceptability* of a conditional is governed by the conditional probability of its consequent given the antecedent. Although the Adams Thesis has been widely accepted as self-evident, it turned out not to hold as a general rule. At best, one can argue that the acceptability of a conditional correlates with the corresponding conditional probability. However, by classifying conditionals depending on the type of an inference they express—following the philosophical tradition of classifying inferences as deductive, inductive, and abductive—Douven and Verbrugge obtained positive results, too. For *deductive inferential conditionals*, the strongest version of the thesis holds: the acceptability of a conditional approximately equals the corresponding conditional probability. For the *abductive inferential conditionals*, a high correlation between the two measures has been observed, while in the case of *inductive inferential conditionals* we can only talk about moderate correlation. Building upon Douven and Verbrugge's study, Krzyżanowska, Wenmackers, and Douven (2013: 'Inferential conditionals and evidentiality,' *Journal of Logic, Language and Information*, 22(3), 315-334) showed that the type of an inferential connection between antecedents and consequents does not only affect the strength of the correlation between the acceptability and conditional probability, but it also affects the way conditionals interact with epistemic modals inserted in their consequents.

While the results due to Douven and Verbrugge highlight the significance of the connection between antecedents and consequents for people's interpretation of conditional sentences, they do not allow us to conclude anything about its status as a semantic or pragmatic aspect of their meaning. More recent results, however, suggest that the connection should at least be considered as belonging to the conventional content of conditionals, if not even its truth-conditional content. Skovgaard-Olsen, Singmann, and Klauer (2016:' The relevance effect and conditionals,' *Cognition*, 150, 26-36) showed that the connection, understood in the probabilistic terms as the antecedent's probability raising effect on the consequent (so called *probabilistic relevance*) affects people's probability ratings. More specifically, Skovgaard-Olsen et al. (2016) investigated the thesis, typically referred to as *The Equation*, that the probability of a conditional equals the corresponding conditional probability. While it is believed to be the most robust finding about indicative conditionals, the Equation turned out not to hold for all conditionals, but only for those whose antecedents are relevant for the consequents.

Another line of empirical research that motivates the semantic approach to the connection between antecedents and consequents does not concern the semantic content of a conditional directly, but it shows that a purely pragmatic treatment of the connection is problematic. For instance, Krzyżanowska, Collins, and Hahn (2017: 'Between a conditional's antecedent and its consequent: Discourse coherence vs. probabilistic relevance,' *Cognition* 164, pp. 199–205) show also that the oddity of missing-link conditionals is not due to the violation of discourse coherence, that is, that the connection between the clauses of a conditional needs to be something stronger than the common topic understood in discourse-coherence-theoretic terms. Furthermore, the forthcoming paper by Skovgaard-Olsen, Collins, Krzyżanowska, Hahn, and Klauer (2019: 'Cancellation, negation, and rejection,' *Cognitive Psychology* 108: 42-71) shows that the connection cannot be a conversational implicature since a speaker attempting to cancel it is judged by the participants as contradicting themselves. The oddity of conditional's without a connection is also not an instance of a presupposition failure, since it does not project under wide scope negation. Moreover, it does appear to belong to the at-issue content. While the possibility that the connection is a conventional implicature is still open, making it a semantic, but not truth-conditional content, a recent work by Douven, Elqayam, Singmann, and van Wijnbergen-Huitink (2018: 'Conditionals and inferential connections: A hypothetical inferential theory,' *Cognitive Psychology* 101, pp. 50-81) provides evidence that the presence and the strength of an inferential connection affects people's truth value judgements, too.

Given the close relationship between conditionals and hypothetical reasoning, these results are not surprising: after all, in the process of hypothetical thinking, people tend to be interested in the consequences of what they suppose that are related to their suppostions, not merely in things that happen to be true when these suppositions hold. How to exactly account for this phenomenon is an exciting research question that we hope to answer. Stay tuned!

Karolina Krzyżanowska
University of Amsterdam / University of St. Andrews

What's Hot in ...

Medieval Reasoning

> "[P]hilosophers' convictions about the eternity of problems or conceptions were as baseless as a young girl's conviction that this year's hats are the only ones that could ever have been worn by a sane woman".

This passage in Collinwood's *An Autobiography* (Oxford 1939, p. 65) has always resonated with me. The thing is, I am not entirely convinced that Collingwood was right, but he might have been onto something – besides women's fashion. As a historian of medieval philosophy (and a casual historian of fashion), my professional identity is an odd beast, like a unicorn or a chimera. Not in the sense that historians of philosophy are mythical monsters – you can find a few of us wandering around departments of philosophy and it doesn't look like we are particularly close to extinction, yet. But in the sense that we have multiple natures: we are historians *and* we are philosophers. On the one hand, the historian within me *knows* that Lady Philosophy has changed a lot over her long life. My inner historian likes to picture her as an old lady who's had a few plastic surgeries too many and has lost a few bits here and there – oftentimes to replace them with more or less eccentric prosthetics, only to occasionally switch them over again, to keep pace with the ever-changing fashions of the day. Or perhaps my inner historian entertains the idea that Philosophy is not quite a lady, but rather a barely sketched vaguely written role interpreted by different actors; or even better an artificial person, like an institution: what that institution is and does changes with the people inhabiting it, its practices, its reformations and, overall, the times, and yet the institution itself is still in some sense the same. Some days, my inner historian thinks of Philosophy as a bit of both – the old lady and the institution – , i.e. the same sort of patchwork creature that we, her historiographers, are. Long story short and out of metaphor, a good chunk of philosophical issues and conceptions, that were essential at some point or another in the past, doesn't count as philosophical at all in our eyes – think, for example, of some of the things historians of ideas, theology, or even science are interested in. The converse would probably be just as true. At the end of the day, my inner historian acknowledges the data and interprets it, trying to tell a coherent story of the hows and whys of this historical development. On the other hand, the philosopher within me is more conflicted, which is not surprising. My inner philosopher wants to believe that philosophical questions and theories, for the most part, are not unsolvable conundrums or unchanging truths – the very same we have been dealing with since the dawn of our discipline – that we have been doomed to address until the end of time, with no real hope of resolution. What a boring and utterly hopeless endeavour would philosophising be then! Yet, my inner philosopher has a recognition that there is some sense in which the stuff she is doing is the same kind of stuff that the philosophers of the past were

doing, i.e. philosophy. That recognition might however be misguiding, or even delusional, not merely because my inner philosopher might be that bad at philosophising, but because the recognition itself comes with a preconception – very much shaped by our time and curricula – of what philosophy is supposed to look like. My inner philosopher, then, wonders about whether there is some deep core to philosophy, i.e. a set of essential features making something into "philosophy", i.e. a common denominator shared by anything that was, is and will be philosophy, across ages and continents. It's a tempting thought, of which my inner philosopher – fancying herself to be as nominalist as they come and being good friends with my inner historian – is pretty weary.

The problem is even more evident as far as logic and its history are concerned.

TBC...

Graziana Ciola
Durham University

Uncertain Reasoning

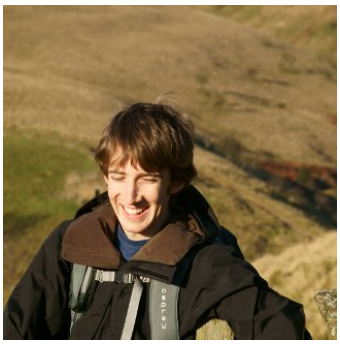

Here are some thoughts that have been inspired by reading Wallmann and Hawthorne (Admissibility Troubles for Bayesian Direct Inference Principles, forthcoming in Erkenntnis), but this isn't really about their paper in particular. The main thought driving this piece is that I think the grounds for thinking that conditionalisation is a general feature of rational updating are quite weak. In particular, I think that arguments for conditionalisation don't provide a strong justification for conditionalising on some kinds of evidence.

Let's look at the arguments for conditionalisation, starting with the betting argument. What this shows is that your betting quotients in events ought to be probabilistic. This argument seems most natural when considering events, or propositions concerning straightforward matters of fact. Betting on propositions whose truth is difficult to determine for whatever reason seems somewhat conceptually confused. For a discussion of betting on vague events, see Milne "Bets and Boundaries" (2008 Studia Logica). The important point is that betting only makes sense under the assumption that the truth values of the propositions bet on can be determined to both parties' satisfaction.

Consider betting on chance events. How does one verify that a chance event is true? The defender of the orthodox broad scope of conditionalisation will say that in fact you don't need to figure out whether the chance event is true, since, if you don't update appropriately there's a set of bets on standard well-behaved propositions such that you consider them acceptable and yet you will have a negative expected gain. Now, at the moment I'm happy to let that argument stand, but note that what it really requires is that learning the chance proposition has the right effect on your credences in other propositions, not that the update is necessarily by conditionalising on a chance proposition.

If we are careful to distinguish update – the general rational change in belief – from conditionalisation – update by moving to the conditional probability – then we can put it like this: betting arguments for conditionalisation only really make sense when the events you bet on are events whose truth can be determined, and betting arguments for the "principal principle" or other direct inference principle target updating, not conditionalisation.

The same kind of pattern can be discerned in other arguments for conditionalisation. Consider the argument that conditionalising minimises the information gain subject to the constraint that the evidence gets probability 1. (See for instance Diaconis and Zabell "Updating Subjective Probability" (1982 Journal of the American Statistical Association)). But this formulation of the argument is lacking. What is really shown is that updating by minimising information gain subject to a constraint is equivalent to conditionalisation only when the only constraint is that the evidence gets probability 1. But updating on a chance proposition obviously imposes a second constraint that the updated credence should respond to: not only should the chance proposition get probability 1, but the proposition the chance proposition is about should have whatever value the chance proposition says it should. Minimising information gain subject to those two constraints needn't coincide with conditionalisation.

This brings us back to Wallmann and Hawthorne's paper. They argue that admissible evidence is a tricky topic. They show that on purely logical grounds, lots of surprisingly innocent propositions can defeat direct inference. But, their arguments rely on demanding that updating on a chance proposition is by conditionalisation. If we deny that, then perhaps there is a route to having a more plausible theory of admissible evidence.

Seamus Bradley
Philosophy, University of Tilburg

Mathematical Philosophy

The concept of ecological rationality has captured the hearts and minds of many psychologists and philosophers (Hertwig & Pedersen, *Minds and Machines*, 2016, 26:1–8). In fact, I started my Ph.D. as a believer. The idea is that the apparently simple decision procedures, "heuristics", which are used by cognitive agents such as ourselves, are *ecologically rational* insofar as they make use of the relevant regularities of the environment in which the decision is usually made. The claim that humans use ecologically rational heuristics is two-fold. It consists of both a normative and a descriptive part.

Strictly speaking, ecological rationality is a normative theory that studies the relationship between a given decision procedure and an environment, and how the valuation of a particular decision, as being rational or not, depends upon it. It is usually

accompanied by the adaptive toolbox theory – a psychological account listing the possible decision procedures at an agent's disposal. As partly suggested by the name, the idea is that there's a multitude of such decision procedures, each of which was acquired through an adaptation to different environmental characteristics.

The descriptive premise of the fast and frugal heuristics research program (FFH) is that humans use such procedures and that they mostly do so in an ecologically rational way. It's one of the many offsprings of Herbet Simon's concept of bounded rationality (Simon, *Annu. Rev. Psy.*, 1990, 41:1–19.). What sets FFH apart from the others, for instance the "heuristics and biases" program (Kahneman & Tversky, *Science*, 1974, 185:1124–1131), is the emphasis on the environmental component, rather than scarcity of cognitive resources. The idea is to postulate some general cognitive limitations, *e.g.* a working memory capacity, together with the constraints imposed by the environment, *e.g.* a cost of obtaining information, and derive some testable predictions about human behavior.

Heuristics are said to be frugal in the sense that they ignore information, and fast in the sense that frugality leads to less computations. The more recent theoretical developments in the literature downplay the importance of frugality in the face of a cognitive load. Instead they emphasize the importance of the robustness of heuristics that has evolved as a response to opaque environmental uncertainty, as opposed to a quantifiable risk (Mousavi & Gigerenzer, *Homo Oecon*, 2017, 34:361–379). Ignoring information helps against overfitting. However, it remains assumed that heuristics are "simple" in the sense that they require little computation to be performed.

I've become increasingly suspicious about this descriptive claim. Mostly so only after noticing a rarely discussed detail in a pet quote of proponents of ecological rationality:

> Human rational behavior (and the rational behavior of all physical symbol systems) is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor. (Simon, *ibid.*)

What does it mean for humans to be physical symbol systems? The notion of heuristic together with the idea of bounded rationality is deeply rooted in the research on symbolic artificial intelligence, or classical AI. FFH inherited at least some of the methodology from the research tradition that tried to understand cognitive processes by analyzing computer programs designed to replicate the related behavior. I find this legacy to be very problematic.

The behavior is analyzed in the same way as one would analyze a computer algorithm. The standard way to infer the complexity of a decision procedure, *e.g.*, a heuristic, is to count the number of elementary operations (like compare, recall, increase counter, *etc.*). Think of the Random Access Machine model without assuming that memory calls are costless. The problem with this approach is that it is not really clear what the "actual instruction set" might possibly look like.

Examining our own hand and finger movements goes a long way towards showing that our intuitions about simplicity can be completely off when we move from algorithms to functioning of a neurobiological system. Moving a single digit requires more control and is metabolically more costly than a "composed" movement of forming a fist, or grasping an object (Schieber, *TINS*, 1990, 13:440–45).

I write about this in great length elsewhere (under review). There I argue that the understanding of human behavior and cognitive processes underlying FFH is at odds with recent empirical findings and our best available theories of how decision processes are carried out in our brains. Going a step further, I claim that the role of simple cognitive algorithms can be replaced by the notion of simple neural mechanisms, while preserving, and even further substantiating the intuition that our behavior is ecologically rational.

Most, if not all, tangible definitions of simplicity (think of complexity theory) assume a particular type of hardware, which is just to say a *type of computer*. It seems reasonable to ask what *type* of computation is actually being performed. Currently I'm looking at different assumptions about the type of computation that is postulated by FFH and how this relates to Simon's original account of the bounded rationality.

There is a lot to be said about the topic and this is where my interests lie right now. Shoot me an email at `gasper.stukelj@lrz.uni-muenchen.de` – I'd be happy to talk about it!

Gašper Štukelj
Munich Centre for Mathematical Philosophy

Evidence-Based Medicine

Almost since EBM's conception, questioning the core tenets of the movement has been a part of the philosophy of medicine. In recent years, the focus has shifted from criticising, to making the case for expanding, what counts as the 'best evidence'. At the tail end of 2018, a letter in the BMJ, authored by philosophers of medicine, and undersigned by 42 clinicians and philosophers, urges EBM to "acknowledge the importance of understanding causal mechanisms". The letter identifies that evidence in EBM, is evidence of causation. Typically, evidence of the mechanisms that explain how an intervention works is not used to contribute to establishing causation. The signatories think this is a mis-guided approach, and make this case by appealing to recent work in the philosophy of medicine (see the list of references in the letter for key work in this area, and also the work of EBM+). Conclusions from this work are summarised in 8 key points. Implications for the practice of EBM that these conclusions make are deep and for many will be controversial. Instead of just explaining how and intervention works, causal mechanisms can be used to justify *whether* an intervention works. Evidence should not just come from statistical sources, but should be obtained from laboratory studies, non-randomized clinical trials, and patient accounts of their own diseases. The evidence of causation these sources provide should not privilege causal interactions at the physiological and molecular level, but should be extended to higher-level factors such as the individual or social. This expansion of the notions of evidence and causation run counter to many accepted ideas in EBM. However, the changes they suggest are claimed to only improve on EBM, not supplant it. The philosophical literature supports this view, and support in the clinical community is growing.

However, a recent comment in prominent medical journal *The Lancet* shows how the dominant view is alive and well. The aim of the comment is at calls for randomised controlled trials to be 'replaced' by 'real-world trials'. This amounts to replacing randomised studies with non-randomised observational

studies. A driving reason for this approach is that such real-world trials can use larger data sets that are more representative of the population the intervention will eventually be targeted at. RCTs are well known to be more restrictive in both the size and characteristics of the sample population, which makes extending results to a target population difficult. On the other hand, observational studies are argued to be inherently at greater risk of systematic bias. Statistical techniques that account for bias are the proposed solution. The comment argues in turn that this is no replacement for randomisation, because of the "power" of randomisation to balance for all confounders known and unknown and "ensure that the true effect of an intervention can be assessed".

This is an argument that has been treated at length in the philosophical literature, and the special power of randomisation is hotly contested. One might think then that this comment is misguided in premising its whole argument on this point. Of course, proponents of randomisation will argue in turn that critics are just wrong, and randomisation does do what it is supposed to. However, we can avoid this issue and still reject the main thrust of the comment's argument. This is because the authors take issue with arguments that RCTs should be replaced by other sources of evidence. As the letter in the BMJ should indicate, a more reasonable approach is to supplement RCTs with other sources of evidence. Instead of maintaining that RCTs provide the best evidence in all cases, a more nuanced approach to evidence assessment may allow us to identify times when a randomized study alone is appropriate, or when we also need to obtain other kinds of evidence of causation. It is also the case that sometimes the sort of large scale RCT that *The Lancet* comment authors argue is the best source of evidence is not feasible to perform. In that case we are forced to look to other sources of evidence. As the authors of the BMJ letter stress, this is a way to improve medicine rather than to revise it. More generally, the old and tired argument over the primacy of randomisation should be put to bed, and we should focus instead on working out how to establish causation without relying solely on one method.

D.J. Auker-Howlett
Philosophy, Kent

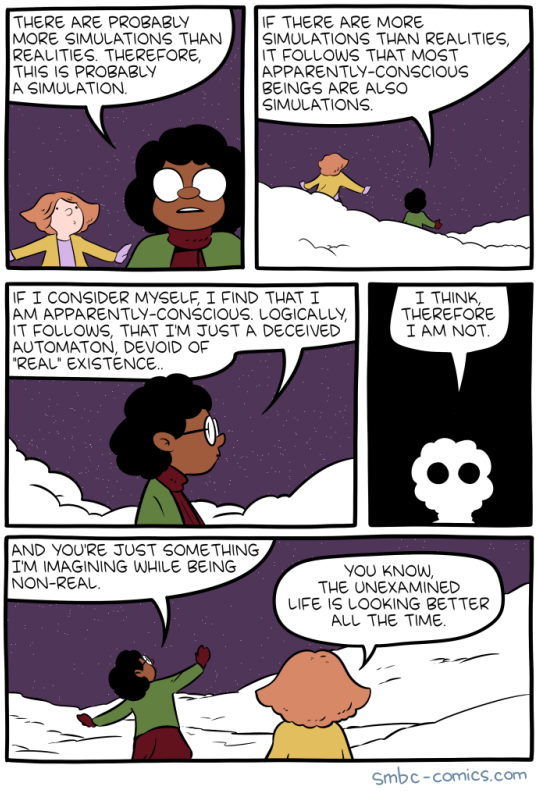

Events

February

SURe: Scientific Understanding and Representation, Bordeaux France, 5–6 February.
DMaMG: Dark Matter and Modified Gravity Conference, Aachen, Germany, 6–8 February.
Con&Ex: Concepts and Explanation Conference, Dusseldorf, 7–8 February.

March

PhiMet: Workshop on Philosophical Methodology, Barcelona, 14–15 March.
ArgMap: Argument Mapping, Nova University of Lisbon, 15–18 March.
BCF: Beyond Curve Fitting: Causation, Counterfactuals, and Imagination-based AI, Stanford, California, 25–27 March.
PTS3: Proof-Theoretic Semantics: Assessment and Future Perspectives, Tbingen, Germany, 27–30 March.
M-S PoS: Mid-South Philosophy of Science Network, Lexington, Kentucky, 29–30 March.

April

ResLog: Reasoning, Argumentation and Logic in Natural Language: Experiments and Models, Ruhr University Bochum, 3–5 April.
Formal Methods and Science in Philosophy III, Dubrovnik, Croatia: 11–13 April,

.

Courses and Programmes

Courses

SSA: Summer School on Argumentation: Computational and Linguistic Perspectives on Argumentation, Warsaw, Poland, 6–10 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.
LogiCS: Joint doctoral program on Logical Methods in Computer Science, TU Wien, TU Graz, and JKU Linz, Austria.

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

Postdoc: in Theoretical Aspects of Graphical Models, Pompeu Fabra University, Barcelona, open until filled.
Associate Professorship: in Medical Philosophy, University of Aarhus, deadline 4 February.
Postdoc: in Emergence in the Natural Sciences, Lisbon, deadline 7 February.
Lecturer: in Statistical Science, University College London, deadline 9 February.
Postdoc: in Evidence Synthesis, Paris Descartes University, deadline 10 February.
Postdoc: in Philosophy and Sociology of Science, Berlin, deadline 15 February.
Postdoc: in Statistical Standards and Evidence Amalgamation, Polytech University of the Marche, deadline 15 February.

Studentships

PhD: in Causal Inference, University of Sheffield, deadline 1 March.
PhD position: in Philosophy of Science, University of Sofia, deadline 31 March.