The Reasoner

Contents

Guest Editorial

The rapid spread of (mis)information by modern communication is frequently highlighted as a characteristic of a 'post-truth society'. The general public often lacks scientific and media literacy, which makes it easier for misleading beliefs to spread. Formal modeling has been a popular way of addressing the spread of false beliefs or divisions between people within networks of agents.

I was very happy to interview Cailin O'Connor, Associate Professor of Logic and Philosophy of Science at the University of California, Irvine who is working on precisely these topics. We talked about her new book *The Origins of Unfairness: Social Categories and Cultural Evolution* that is coming out this month, and in which Cailin analyzed the origins of inequality across different societies. In the interview, she also addressed the spread of misinformation and the role of formal modeling in social epistemology.

Vlasta Sikimić
University of Belgrade

Features

Interview with Cailin O'Connor

Vlasta Sikimić: Thank you very much for giving an interview for the Reasoner. Your work is on formal modeling in philosophy, more precisely on formal models in social epistemology. Could you tell us more about your background and how you decided to study philosophy in the first place?

Cailin O'Connor: I had a really unusual journey to philosophy. As an undergraduate I did not study philosophy at all. I started in biology but then actually got my degree in filmmaking and I was working in documentary film for a few years. Then I decided to get a PhD and I was thinking about combining interests in the sciences and humanities. It was around then that I found out about philosophy of science. With no background in philosophy I started hanging around the logic and philosophy of science department at the UC Irvine. I sat in some classes. It led me to write for them and using that I ended up being able to put together an application package for graduate school in philosophy. I started graduate school without having formally studying it, but having sat in on these classes in philosophy.

VS: So, it is similar to Feyerabend's idea that chance plays an important role in scientific careers?

COC: It was by a complete chance and the only reason why I

found about the field was that my then boyfriend, now husband was studying philosophy of science. It was the first field I heard of where you really can combine humanities with the sciences in a strong way.

Getting into modelling was pretty chancy, too. I had originally thought that I would study the neurobiology behind aesthetics. I was interested in combining film and biology, for instance what makes us like different images or landscapes from the neurobiological view. Instead what happened is that I just started taking some classes on modelling, game theory, and evolutionary game theory, looking at some work done in logic and philosophy of science and getting into it.

VS: Among other things you worked on polarization of social groups. In your experience, how should we reason with people who have extreme views?

COC: It's been observed many times that just presenting people with extreme views with evidence will not solve the problem and rehabilitate their false or extreme beliefs. For example, in the US we have huge problems with people who don't vaccinate their children and believe that vaccines are really dangerous. In fact, lots of countries have that problem. People tried to introduce anti-vaxxers to the evidence that vaccines are safe and change their mind. There is lots of scientific information that vaccines are safe, but that does not convince them. One thing I talked about in my book with James Owen Weatherall, *The Misinformation Age*, is how much trust is important for changing people's mind. People decide who to trust and then they listen to the people they trust, but they don't necessarily listen to other people. An issue in the vaccine case is that people create communities of anti-vaxxers and only trust the evidence shared in their community, but don't trust the evidence outside of their community. So, the question is how could you put the important information into the hands of someone who these people with extreme views do trust. The idea is to ask someone with whom they share a lot of views and other beliefs and then try to convince them. They need someone who can tell them "I understand where you are coming from and I really get who you are, but here is why you are wrong". Something like that would help in these cases.

VS: You just mentioned your recent co-authored book *The Misinformation Age: How False Beliefs Spread*, which tackles a socially and politically relevant topic. In your opinion, how should we intervene and more generally act in what is popularly called 'post-truth society'?

COC: There are so many things that we can do. False beliefs have been an issue in all human societies. One of the things that is important to understand is how much knowledge people get from other people. We are really social learners. A lot of things we believe in is just spreading from person to person. Just think about why do you believe that Earth goes around the Sun and not *vice versa*. You probably didn't do any calculation to show that's the case. People told you and then you believed them. Whenever you have a system like that, it will be the case that false beliefs spread too. In this book, we argue that if we think about fixing false beliefs and helping our 'post-truth society', we need to really think about the social learning aspect of humans. To give an example, we talked about anti-vaccination and we want to understand how to make people have a better belief. We need to think what are their social realms and why do they have their original beliefs. Often, they are surrounded by other people who have such beliefs. Often, they don't trust people who have different beliefs. So, they will not

trust the better evidence coming from a wider community. To fix our problem we need to think about trust and we need to think about community structure, whether there is some way to change that structure by introducing people with better beliefs to the community, etc. We have to know a lot about how peoples' minds work. We can intervene in many ways, ranging from changes in social media to promoting individual scientific and media literacy, to government interventions.

VS: How is formal modeling helping in solving these problems?

COC: Formal modeling is helpful in understanding how the knowledge works in order to intervene on it. The reason why formal modeling is particularly useful is that beliefs are often created and spread in the super complex system that's extended in time and space. Sometimes, it can be really hard to grasp who originated some false belief, who did they first share it with, why did those people trust them, why did they share it with another person, and how did it spread throughout our whole network of people. You can use a model to explore how this process can happen in an artificial scenario with simplified individuals and rules of interaction. Once you build these models you can isolate different things and ask why do people like to conform with each other, how does that particular psychological factor influence spread of beliefs and then to start understanding these really complex processes.

VS: In this light, where do you see the room for philosophers in the modern society? Clearly there is the importance of public presence of philosophers and you are personally rather active in this respect.

COC: I think there are a lot of important things philosopher can do in modern society. I'm going to focus on philosophy of science. One thing we see which is an issue is a lack of scientific literacy in public. Often people think that one study either shows that something is true or false. This is a naïve view of science. That is not really how science works. It often takes dozens of studies and a lot of debate within a community to settle a particular question, because scientific evidence is probabilistic. For example, if you look at 1950s, you can find dozens of studies that show there is no correlation between the cigarette smoking and cancer. The reason is that not everyone who has cancer smokes and not everyone who smokes has cancer. You are going to find these studies that show the wrong thing. But if you look all the studies they did at the time, there is overwhelming evidence that tobacco does cause cancer. That's why having this naïve view of science can be credible misleading. Journalists would see one study and publish the popular article based on the one study. And everyone would think "Oh now we know X", but they don't, because that is not how scientific evidence works. I think that philosophers of science can play the role in helping people understand what scientific evidence is, how it works and intervening in cases where people are understanding evidence incorrectly.

VS: Together with your group, you also worked on empirical approaches in philosophy. What you is the potential of experimental philosophy in your opinion? And how does it connect with formal modeling?

COC: Sometimes people think that the only proper methods of philosophy are thinking, reasoning, introspection, and maybe logic. Yet, sciences have developed all these fantastic tools for finding out about the world. Philosophers should not work their hands tied behind their back, but use all those fantastic methods that are out there. What I think is special about

our discipline is the kind of questions we are interested in, but not necessarily the kind of methods that we use. Let's just use the best methods there are. Experimental philosophy is using good experimental methods from different disciplines like psychology, and increasingly economics, to answer philosophical questions and that seems really important. When it comes to formal modelling, one important thing that simple models do, is direct our attention to things that might be happening. We are not 100% sure they are happening in the world, but if they are happening in a model there is a good chance that they are happening in the world. That can direct us to take up empirical studies to test something we would not test otherwise. Here is an example, Justin Bruner discovered the effect in evolutionary models that occurs when you have two groups of agents interacting and you have a minority group and majority group and they're bargaining. Sometimes a minority group ends up getting less in the bargain just by being in a small group. The reason is that they are meeting their outgroup so often that they learn quickly to accommodate them and then the outgroup slowly learns to take an advantage of that. Justin and I called this 'the cultural red king effect'. We found this effect in many different models, which means that this a very robust modeling effect. Recently, Hannah Rubin, Aydin Mohseni and I decided to do a study this effect experimentally by having humans in a lab and letting them bargain to see does the small group end up getting less and we found they did. Without models we would have never designed that experiment. So, there can be these nice connections between what we find in the models and empirical or experimental philosophy.

VS: Your new book should appear this month [in July]. It's called *The Origins of Unfairness: Social Categories and Cultural Evolution*. Could you tell us what to expect?

COC: That is related to this project on the cultural red king effect that I was just talking about. This book looks at cultural evolutionary models using game theory. The biggest question in the book is "Why do we have unfair or inequitable norms and conventions in society?" Basically, every society has unfairness especially unfairness over social categories. Social categories are for example genders, different racial categories, religious categories, age categories, etc. When we have categories in society, often it's the case that when people divide resources or do work, one group gets less and the other gets more. My big question was why would this be such a persistent and stable pattern. Basically, as soon as u get groups you get unfairness. I did a lot of modeling. In the modeling work you have population of agents that are divided into two groups. These groups are basically identical, I treat them all the same. The groups bargain with each other and you can show in these basic models that it very commonly happens that one group gets more and the other one less. Just the fact that there are groups creates asymmetry. From this starting point, I expand the exploration and look at, for example, why do we have gender division of labor in society. Men and women have different jobs in almost every single measured society. Why is it the case that women tend to be disempowered in different societies? How can we explore the minority and majority group? Is the minority disadvantage a common thing? I looked at network models to ask if you have unfairness or inequality does that lead to people breaking away from the group that is treating them badly and only be with their ingroup. I also used these models to ask: "Is there a way to intervene to the unfair norms and conventions?"

VS: And how should we intervene?

COC: When we think about unfairness in a cultural system, we need to think about the underlying strategic situations. It is often that people are engaged in the situations where everybody to some degree wants the same things, but to some degree wants something else. There is some common interest and some conflict of interest between the groups. What this means is that if some group is disadvantaged, its members often have a way to change the behaviour of the other group because there is some common interest and because everybody wants the bargain to be made. If you think about household division of labor, women do more labor in general, household labor, market labor, and they tend to have less free time. Women might want that to be different, but that is a norm, and everybody has learned that pattern from their parents and passing it on their children. This pattern is hard to change. However, there is a common interest between women and men since everyone wants that the household runs well. If women want to make men to work more they should say: "I am just not going to work anymore until you pick up your fair share." The common interest that men have in a household can push them towards a new type of behaviour. So, I talk about strategic scenarios that can get us a more equal division.

VS: Finally, would you like to share some experiences of how it is to be a woman in formal philosophy?

COC: My area of philosophy – philosophy of science – is strongly male dominated. A few years ago, when we look at the number of submission to the Philosophy of Science journal, only 16% of them were women which gives some notion of how many women are philosophers of science. It is not a lot. Logic is probably even worse, it is also very male dominated. What I do, because it is intersection of applied mathematics and philosophy, has also very few women. I had a very good experience *vis-à-vis* and generally not felt that I was treated badly by the men in the profession. Some people have had problems, but I mostly had nice interactions. Still, there have been some structural issues that have been hard. In my third year of grad school I gave birth to twins. Finishing my dissertation with twin infants was demanding. I was breast feeding and they are so physically dependant on you. A lot of people in academia do not understand the physical toll that pregnancy, breast feeding and infant care takes on you. It was very challenging. I had issues at conferences where I would need to pump breast milk and there were nowhere to go and nowhere to store it and nobody has done anything about it. Just being at a place that has been mostly created by men, even if they are doing a good job and being good people, there are sometimes not enough accommodations for women's bodies and women's needs.

News

Calls for Papers

What's Hot in . . .

Medieval Reasoning

Apparently, it's summer. And, with summer, comes the annual resolution of catching up with our writing and taking a few days off to go somewhere ideally warm and bright with our friends and family – making sure to carry the above-mentioned writing-to-be in our hiking backpacks and beach bags, along with maybe a puzzle magazine or two: crosswords, sudokus, cryptograms. . . name your poison. Summer is the time to indulge. But if you are on the market for something new, you could try your hand at medieval sophisms (*sophismata*). A *sophisma* is not a clever yet false argument, nor has it much to do with its ancient homonym; it is instead a sentence that, carrying some logical or grammatical complications, needs to be disambiguate, usually in some non-obvious way. While already discussed in 12th century treatises, *sophismata*-literature flourishes in the 13th and 14th centuries; at that point these sophisms had become an essential feature of both pedagogy and philosophical argumentation. But, most importantly, they make for delightful little puzzles. For an exhaustive general overview of medieval sophisms, take a look at Fabienne Pironet's and Joke Spruyt's entry in SEP, which includes a rich bibliography and links to online resources if you are curious to know more. For now, I leave you with a few medieval puzzles by John Buridan (from Gyula Klima's translation of the *Summulae de Dialectica*, 2001) to solve on the beach or to ponder over while on your hiking trip.

Ch.1, sophism 6: "Nobody lies"
"Proof: to lie [*mentire*] is to go against one's mind [*contra mentem ire*], i.e., to assert the contrary of what one has in mind. But this is impossible, because to every spoken proposition there must correspond a similar one in the mind. For one would not assert that which one did not understand; therefore, etc. The opposite is obvious."

Ch. 2, sophism 3: "A chimera is a chimera"
"A chimera is a chimera "This is proved by [the authority of] Boethius, who says: 'No proposition is truer than the one in which the same thing is predicated of itself.' Again, if it were not true, then this would only be because the term 'chimera' supposits for nothing. But this does not prevent its truth, for the [proposition] 'Aristotle's horse walked' is true, even though the subject supposits for nothing. This I prove: for that subject does not supposit for something other than Aristotle's horse; and that is nothing [. . .]. The opposite holds because this is an affirmative proposition and its terms supposit for nothing, and such affirmative propositions are false."

Ch. 4, sophism 4: "An old man will be a boy"
"Proof: this is equivalent to 'He who is or will be an old man will be a boy', and this is true for Antichrist. The opposite is argued by induction: this old man will not be a boy, nor this, nor this, and so on for [each] single one — for this is true whichever old man is pointed out — therefore, no old man will be a boy. Again, one can argue just as was argued in the

preceding sophism, concerning possibility as well as the future.

Graziana Ciola
Durham University

Uncertain Reasoning

One interesting and fruitful way to think about how your rational beliefs are constrained is to think about what sort of invariance principles they satisfy. For example, if you endorse some form of the "principle of indifference" or the "principle of insufficient reason" then you think that if you have no evidence regarding two propositions, you ought to believe them to the same degree. Or (even if you don't buy the principle of indifference) different tosses of the same coin should affect your rational beliefs in the same way. In short, if some precondition as regards some propositions holds then your beliefs should be invariant under swapping around those propositions. That's a pretty vague, sketchy way of describing the kinds of things I have in mind, but some of that indeterminacy is deliberate, since I'm talking about a pretty broad class of principles.
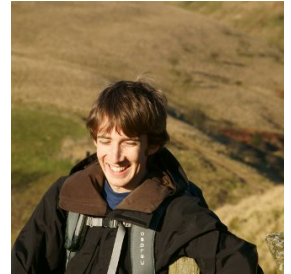
Some forms of this "invariance reasoning" are frutiful and productive ways of building rich and interesting theories of rational belief in some domain. One classic example is de Finetti's work on "exchangeability" (de Finetti, Theory of Probability, Wiley 1974). The basic idea is that it doesn't matter what order you learn about independent tosses of a coin; whether you learn about toss 1 first and then toss 2, or toss 2 first and then toss 1, you end up with the same final belief. If this is so, then we can represent your subjective probability in a compact and reasonable way as a mixture of views about the coin's bias. This work has inspired a quantum generalisation of the approach (e.g. Caves, Fuchs,and Schack. "Unknown quantum states: the quantum de Finetti representation." Journal of Mathematical Physics 43.9 (2002): 4537-4559.); and a wealth of interesting work in the Imprecise Probabilities framework (e.g. Chapter 10 of Troffaes and de Cooman "Lower Previsions", Wiley 2014).

Another classic example of this approach is Carnap's Continuum of Inductive Methods, where your probability is invariant under various kinds of transformations. That is, Carnap's probability functions are arrived at by presuming that certain sentences in a predicate language should be given the same probability. The idea is that the only discriminations that can be made in the absence of evidence are on the basis of the logical structure of the sentences involved. The "Pure Inductive Logic" project of Paris and Vencovská (Cambridge University Press 2015) extends and generalises both the de Finetti and Carnap approaches to invariance. They show that Carnap's favoured probability functions are characterised by the invariances they satisfy.

Others have reasoned that your epistemic state ought to satisfy some sort of invariance principle, and used that as an argument against a particular formal model of belief. For example, John Norton has argued that Bayesianism cannot provide the right theory of induction since it does not adequately represent ignorance (Norton "Probability disassembled." The British

Journal for the Philosophy of Science 58.2 (2007): 141-171.; Norton "Disbelief as the Dual of Belief." International Studies in the Philosophy of Science 21.3 (2007): 231-252.). His argument for this rests on arguing that ignorance requires certain kinds of symmetry in our attitudes that cannot be accommodated within a standard Bayesian picture.

One particularly interesting kind of invariance is a sort of "invariance under choice of language", or choice of state space. Many theories of logical probability or objective Bayesianism fall foul of these kinds of principles. Imagine you're trying to figure out what to think about the probability that the next car you see is red. You could describe the state space as "red/not red" or you could use the more refined state space "red/orange/yellow/green/blue/indigo/violet". A naive application of the principle of indifference would yield different probabilities for red in these two cases. Arguably we don't want this sort of dependence on the structure of the state space when figuring out our probabilities. Peter Walley makes an argument along these lines in order to justify the use of Imprecise Probabilities (Walley, "Statistical Reasoning with Imprecise Probabilities", Chapman and Hall, 1991). Norton also appeals to cases like the above (Norton "Ignorance and indifference." Philosophy of Science 75.1 (2008): 45-68.). One might respond that the language you use encodes something about your evidence, and so it's undesirable to wash out that information by making your probability "too" invariant.

Perhaps a stronger argument against too much invariance under choice of language points to how hard it is to have a non-trivial epistemic state in those circumstances. Walley shows that the vacuous prior (the credal set containing all probability functions over the state space) has the property of being invariant under choice of language, but it is totally uninformative! Jon Williamson argues that to require a probabilistic logic to be language invariant is to trivialise the entailment relation ("From Bayesian epistemology to inductive logic." Journal of Applied Logic 11.4 (2013): 468-486.). Williamson's version of language independence might be too strong, and that limits the force of his claims, however, Halpern and Koller provide a more careful analysis of language independence and show that any inference procedure that is language independent is "essentially entailment", which means that there are no interesting inductive inferences that can be made in a language independent setting ("Representation dependence in probabilistic inference." Journal of Artificial Intelligence Research 21 (2004): 319-356.). Finally, Benjamin Eva shows that if you take an ordinal "comparative confidence" approach to rational belief, if you want your confidence relation to be invariant across languages, then your confidence ordering captures just the logical structure of the propositions and nothing more (that is, it is isomorphic to the Lindenbaum algebra of the language)(Eva, "Principles of Indifference" The Journal of Philosophy forthcoming http://philsci-archive.pitt.edu/16041/).

So, here's a research question for you: is there an interesting middle ground? A moderate language invariance constraint that is plausible, fruitful but that doesn't trivialise the resulting epistemic state?

Seamus Bradley
Philosophy, University of Leeds

Mathematical Philosophy

The question of which physical systems perform computations, or which computations they perform, might seem quite trivial, but it is in fact very difficult to answer. If we were only concerned with artifactual computers, i.e. physical computers that we have constructed ourselves, then this might not be such a problem, but the issue becomes more pressing if we are interested in the computational status of natural systems, such as the human brain. A foundational assumption of much of cognitive neuroscience is that the brain is a computer, and that thinking is a kind of computation, and yet it is not always clear what this claim really amounts to, or how we might go about testing it. Philosophers have offered several different accounts of computation in physical systems, each of which suffers from some well known problems.

Simple mapping accounts say that a system performs a computation if we can identify a mapping between its physical structure and the abstract mathematical structure of the computation. The problem is that (without any further constraints) it is almost trivially easy to identify a mapping between any physical structure and any mathematical computation, such that it will turn out that every system performs every computation (Sprevak, M. 2018. Triviality arguments against computational implementation. *Routledge Handbook of the Computational Mind*). Most researchers have wanted to avoid so-called "unlimited pancomputationalism" of this kind, which makes it unclear why we should be interested in computational explanations, so they introduce constraints on the kinds of mapping that are considered legitimate. A common constraint is to say that the causal structure of the physical system must match up with the abstract structure of the mathematical computation, where every transformation or operation in the mathematical structure corresponds to a causal interaction in the physical structure (Chalmers, D. 1994. On implementing a computation. *Minds and Machines* 4/4: 391-402). This might still result in a form of "limited pancomputationalism", where every physical system performs at least one computation, but for many this is preferable to unlimited pancomputationalism.

Semantic accounts say that, in addition to possessing the right kind of causal structure, a physical system must also possess representational content if it is to qualify as computational (Sprevak, M. 2010. Computation, individuation, and the received view on representation. *Studies in History and Philosophy of Science Part A* 41/3: 260-70). This means that if a system is performing a certain mathematical computation, its physical states should represent the abstract variables of that computation, or if a system is performing a computation that is about something in the world, then its states should represent whatever it is supposed to be about. This kind of account is particularly popular in cognitive neuroscience, as if we are trying to explain (for example) the perceptual capacities of a cognitive system, then it makes sense to think of that system as representing what it perceives. However, there is a lot of disagreement about what it means for a physical system to rep-

resent something, so it is not clear that this criteria could serve as an objective means of determining whether or not some system (like the brain) is actually a computer. Furthermore, many popular accounts of mental representation already rely on some notion of computation, threatening this approach with circularity.

Mechanistic accounts closely resemble the causal mapping account, but specify that the kinds of causal interaction that qualify as computational are those that have the function of computing, relative to some target phenomenon that we wish to explain (Piccinini, G. 2007. Computing Mechanisms. *Philosophy of Science* 74/4: 501-26). So if we are interested in giving a computational explanation of the visual system, then we must identify parts (components) of that system whose function is to carry out computations. Consequently mechanists must also provide some account of what it means for a physical system to have a function, and how we decide which function each component of a system performs. The most popular way to do this is by appealing to the selection history of the system, i.e. how it came to be the way that it is, either in terms of biological evolution, purposeful design, or learning within the lifetime of one organism (Garson, J. 2017. A Generalized Selected Effects Theory of Function. *Philosophy of Science* 85/3: 523-43). So we might say that the function of neurons in the early visual system is to perform a certain kind of computation if they evolved to perform that computation, or were otherwise selected to do so. As with semantic content, there is no real consensus about how to determine what the function of a system is, which might threaten to undermine the mechanistic approach to physical computation.

One final option is to simply accept the pancomputational implications of the simple mapping account, and to find some other way to make sense of why computational explanations are interesting (in a world where every physical system is computational). This can be done by adopting a perspectival approach to computational explanation, i.e. by saying that even if every physical system computes, not all of these systems are going to be interesting to us (Schweizer, P. 2019. Computation in Physical Systems: A Normative Mapping Account. In Berkich & d'Alfonso (ed.), *On the Cognitive, Ethical, and Scientific Dimensions of Artificial Intelligence*. Springer). The trick here is to find some principled way to distinguish between the interesting and uninteresting computations, in order to avoid the accusation that computational explanations are arbitrary or completely observer-relative. A recent proposal by Tyler Millhouse is that we should do this in terms of the algorithmic complexity of a computational mapping, such that the simpler a mapping is, the more "legitimate" or interesting it becomes (2019. A Simplicity Criterion for Physical Computation. *The British Journal for the Philosophy of Science* 70/1: 153-78). This seems to sit well with most of our intuitions about which systems compute and which don't, and thus avoids the full weight of triviality arguments, although it might not be able to provide a completely objective criterion, as determining the complexity of a mapping still depends to some extent on observer-relative choices about how to carve up the system and which abstract machine to use as a reference point.

As with many philosophical debates, the question of physical computation may ultimately be a matter of deciding which is the least-bad option: accepting that every physical system computes, admitting semantic content or mechanistic functions into our ontology, or adopting a perspectival approach to computational explanation. My own preference is currently for the latter option, but I don't think this is an easy or straightforward choice to make, and I hope that we will continue to make progress on the question of physical computation in the future.

JOE DEWHURST
Munich Centre for Mathematical Philosophy

Evidence-Based Medicine

Critiques of the methods of EBM are not new within the philosophy of medicine. 'Medical methodology' is a field within philosophy of medicine that recently has seen a lot of work building on these critiques to offer suggestions for the way we practice medicine. As an example, here at Kent, we have a 'Health Methodology group'. This year we have held two workshops focusing on recent work in medical methodology. Topics of talks included: the nature of evidence; consensus conferences; placebos; medical negligence; evidence assessment procedures; immunisation programmes. There is obviously a lot of points of focus within this ever-growing field. As we reach our summer break at The Reasoner, it is worth looking at some major and important contributions to studies in medical methodology from the past year.

The past year has stood out to me for the flurry of books that take as their central concerns the methods of medicine. One book that has generated a lot of interest both within and outside of philosophy is Jacob Stegenga's 'Medical Nihilism'. The medical nihilism thesis claims we ought to hold low confidence in the effectiveness of medical interventions. This is a general stance about all (or most) medical interventions rather than this or that particular intervention. Stegenga supports this conclusion by appealing to the high frequency of past failed interventions, the low effectiveness of those interventions that haven't failed, and importantly the *malleability* of the methods of medical research. The best methods of EBM, so the argument goes, are malleable because they admit too much room for subjective choice in their design and application, and in the analysis and interpretation of their results. Financial interests in medicine only exacerbate this problem. This malleability means that evidence of an intervention's effectiveness should not lead us to conclude that it is in fact effective. Importantly, these features of medical methods make medicine irredeemable. No amount of fine-tuning our methods, or 'detail tweaking' as Stegenga terms it, can revive the fortunes of modern medicine. Instead, we should move towards a 'gentle medicine' approach, where we do not aim to treat through pharmaceuticals, but through social and lifestyle improvements to people's lives. Alex Broadbent, in his 'Philosophy of Medicine', is also very sceptical that medicine can cure disease. He ultimately rejects medical nihilism as the thesis that medicine is 'impotent'. He agrees with the nihilist that medicine promises cure but very often does not achieve this goal, but contends that this is not a reason to reject medicine. Instead we should view the point of medicine as concerning inquiry into the nature and causes of diseases, which it does progress in. This makes Broadbent a medical *inquirist*. Any success of medicine in terms of curing will be incidental to this main business of inquiry. He is also very critical of the methods of medicine. Taken together, his claims about both the methods and success of medicine may be why he does not

suggest improving medical *methods* for testing effectiveness.

The EBM+ group are also critical of the current methods of medicine. They think the kinds of evidence relevant to establishing interventions should be expanded to include not only evidence of correlation, typically obtained from traditional EBM methods, but also evidence of mechanisms, typically obtained from the methods of the basic sciences. Traditional EBM does not consider evidence of mechanisms when making decisions about effectiveness. Much work in the past decade has argued for (or against) the *evidential pluralist* position that we should consider evidence of mechanisms. Some of the group have recently released a book - 'Evaluating evidence of mechanisms in medicine' - the core of which develops guidelines for using evidence of mechanisms to establish the effectiveness of interventions. This book indicates a salient point of difference between EBM+ and the other sceptical camp. Where EBM+ have seen inadequacy in the methods of medicine, they have taken the 'detail-tweaking' path in proposing ways to improve the methods: Inadequacy in the 'best-methods' of EBM leads to proposing the use of evidence of mechanisms; Inadequacy in the methods of the basic sciences leads to the introduction of a framework for assessing those methods. Of course, the nihilist or the inquirist may still think those methods succumb to the same problems as do the traditional methods of EBM. The ability to iteratively improve methods is thus one core point of contention between the two camps. Another recent contribution to this area is Donald Gillies's 'Causality, Probability, and Medicine' (a review of which, by myself and previous 'What's Hot' contributor Michael Wilde, can be found here). Gillies, who is a member of EBM+, is supportive of the view that confirming causality (and so effectiveness) in medicine often requires both evidence of correlation and of mechanism. His development of an action-related theory of causality for medicine does not however restrict him to requiring both in all cases. But he does still come out in favour of the ability of the methods of medicine to confirm causal hypotheses, including those concerning the effectiveness of interventions. This highlights another point of contention between the two camps - the ability (or not) of medical methods to establish the effectiveness of medical interventions.

These recent approaches all build on older critiques of elements of medical methodology to build larger theoretical frameworks that say something about medicine more generally. Two loose camps, both sceptical about current medical methodology, are differentiated by how they view the ability to establish effectiveness of interventions and what changes if any can be made to medicine. Wherever one finds oneself on the many dimensions of this debate, the richness of the work that has been done within just the past year should lead us to agree that this ever-growing field is in robust health. But not without need for more, and hopefully effective, intervention.

<div align="right">

D.J. Auker-Howlett
Philosophy, Kent

</div>

Events

July

PoSE: Perspectives on Scientific Error, LMU Munich, 1–4 July.
A&N: Agency and Norms, Berlin, 5–6 July.

ICASE: International Conference on Applied Sciences and Engineering, Moscow, 5–6 July.
AO&C: Abstract Objects and Circularity, Munich Center for Mathematical Philosophy, 6–7 July.
KatSoS: Kant and the Systematicity of the Sciences, Goethe-University Frankfurt, 8–11 July.
MEASD: International Conference on Mathematical Engineering Applications for Sustainable Development, Dubai, 12–13 July.
TMaL: Truthmakers and Language, University of Oxford, 15 July.
PhiLogi: Workshop on Philosophical Logic, Buenos Aires, 31 July–2 August.

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

Assistant Professor: in Mathematics and Statistics, University of Guelph, open until filled.

Teaching Associate: in Theoretical Philosophy, University of Bristol, deadline 3 July.

Post doc position: in Decision Making Under Uncertainty, University of Oxford, deadline 12 July.

Studentships

PhD position: in Philosophy of Science, University of Geneva, deadline 31 July.