The Reasoner

Volume 12, Number 7
July 2018

thereasoner.org
ISSN 1757-0522

Contents

Editorial

Dear Reasoners,

teaching is (nearly) over and so are marking and all sorts of board meetings. A few trips to conferences and workshops and then the well deserved summer break! This would all sound pretty standard, except many of us interpret summer breaks as the perfect time to do some *real* work or at least some good reading. While The Reasoner cannot help you finishing up the book draft, the community of reasoners can certainly help you packing your luggage with good books. So here is a call for contributions to *The Reasoning Summer Reading List*. Please send your list of up to 5 titles along with a one-line description to features@thereasoner.org. Lists will be published in the August issue of The Reasoner.

Hykel Hosni
Università di Milano

Features

Preferences, Utility, and Rationality: Game Theory in the Lab

Here is a very basic question: can we meaningfully talk about the rationality of a subject involved in an experiment and, if we can, how could we achieve this? To answer this question, first of all we have to be sure that we agree on how two words are used in economics, namely "utility" and "rationality".

Thus, what is the meaning in economics of attaching a certain *utility* to *something*?

To answer this question we address it in the most basic setting conceivable, namely by assuming, without loss of generality, that there is no risk or uncertainty involved in the decision. The idea is that this *something*, e.g., an apple, belongs to given set of alternatives (finite, again just to simplify the exposition and without loss of generality), that we call $F = \{apple, orange, strawberry\}$, on which a decision maker (henceforth, DM) has some preferences, where these preferences are captured by a relation $\succ$ on F, that we call *strict preference*. For example, we assume for a DM under our scrutiny that $apple \succ orange \succ strawberry$: this means that this DM prefers an apple to an orange or a strawberry, and an orange to a strawberry. In this very simplistic case, in order to be able to meaningfully talk about the utility that the DM attaches to an apple, the relation $\succ$ has to satisfy two properties called *asymmetry* (i.e., for elements x, y of F, if $x \succ y$, then it is not the case that $y > x$) and *negative transitivity* (i.e., for elements x, y, z of F, if it is not the case that $x > y$ and it is not the case that $y > z$, then it is not the case that $x > z$). If these properties are satisfied, economists say that $>$ is a *preference relation* over F and it is possible to represent $>$ with a function $u : F \rightarrow \mathbb{R}$, called the *utility function* of the DM.

Thus, in our example, since it can be easily verified that these properties are satisfied, we could write that $u(apple) = 3$, or – in plain English – that for the DM an apple is equal to 3 *utils*, and also that $u(orange) = 2$ and $u(strawberry) = 1$: what is crucial, given the original preference relation of our DM, is that $u(apple) > u(orange) > u(strawberry)$ and not the actual numbers we decide to attach to the alternatives (this is the reason behind calling this notion *ordinal* utility: it only captures rankings). It is extremely important to notice that we can also start from our set F and create for a DM a function, called the *choice function* $c : 2^F \rightarrow F$, that for every subset of alternatives gives the element of this subset that the DM would choose (e.g., $c(F) = apple$, $c(\{apple, orange\} = apple$, $c(\{orange, strawberry\}) = orange$, $c(\{apple, strawberry\}) = apple$, and $c(\{x\}) = x$ for every x in F). Interestingly, if this function c satisfies one property, called the *α-condition* (that says that, for any two subsets A, B of F, if A is included in B and $c(B)$ belongs to A, then $c(A) = c(B)$), then we can safely state that the DM whose choices we have observed through her choice function c has a preference relation over F. Hence, her preferences can again be represented via a utility function!

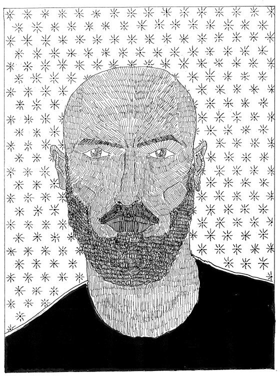

Thus, from the previous paragraph we obtain two crucial points concering modern economic theory:

1. preferences can be *observed* and, in particular, can be *elicited* via the choice function (it is this point that gives empirical content to economics);
2. utility in economics is simply a function that happens to exist for a given DM if and only if the preferences that this DM shows to have satisfy a rational ordering.

Concerning "rationality", the issue is conceptually more problematic. Decision theorists have long fought (and are still fighting) with the idea of providing a formal definition of this notion. For game theorists the problem can be decomposed in two different cases:

a) in the equilibrium-based literature, whose goal is to refine the notion of Nash equilibrium to rule out all those equilibria that are unreasonable *from the point of view of game theorists*, rationality is not formalized, rather it is what arises as the result (outside the model itself) of an equilibrium notion that indeed rules out all the unreasonable equilibria;

b) in the epistemic game theory approach rationality is formally defined from the outset as subjective expected *utility* maximization: a DM is rational if she does not choose any action s that is dominated by another action s' (i.e., such that no matter what the other DMs involved in the game are doing, it is better to choose s' instead of s).

For the sake of our argument, we take the notion of rationality expressed in (b), both because it can be formalized in the language of the theory and because it also intuitively captures the notion that economists seek to capture, but mind that this is (again) without loss of generality! One point needs to be emphasized, namely that a consequence of this definition is that we need a context where a DM expresses preferences enough well-behaved to derive a utility function (otherwise there would not be any utility to talk about in the first place) in order to make inferences concerning her rationality. Hence, rationality cannot be evaluated in a vacuum.

The natural question is now the following: how is all this related to our original problem? Consider an experimenter that sees a subject choosing an action that is dominated from the point of view of the material payoffs provided in the experiment. Can the experimenter infer that this subject is irrational (hence is not playing according to economic theory)? Before moving to an example that should clarify the question, one point. Occasionally in the literature it is mentioned this mythological beast called *homo economicus*, which prototipically is a selfish DM that cares only about her material payoffs. As a matter of fact this description does not capture the behavior of an ideal DM that behaves according to economic theory, since such DM would always maximizes her *subjective expected utility* (which can be different from her material payoff!). In other words, what we wrote about apples and other fruits applies to monetary allocations as well.

The following example will hopefully make more explicit what was written above: take a game between two DMs, call them Ann and Bob, where, given 10\$, Ann has to divide them in whatever way she prefers and Bob has to simply accept her decision: this is called in the literature the *Dictator Game*. One question naturally arises: If Ann divides the amount of money in an equal way, i.e., 5\$/5\$ (where the first number is her material payoff and the second is Bob's), is she kind and rational or selfish and fool? As a matter of fact, we simply do not know, since by taking this single choice, we cannot say anything concerning her utility function and, in turn, we cannot talk about her rationality. To clarify, *if* Ann's utility is a mirror of her (and only her!) material payoff, then we could conclude that she is irrational. Indeed, since her preferences mirror her material payoff, we have that $10\$ > 9\$ > \cdots > 1\$ > 0\$$, and – by assuming that the number of utils coincides with the number of dollars – we have that $10 > 9 > \cdots > 1 > 0$, which in turn implies that, by taking just 5\$, i.e., 5 utils, she is not maximizing her *utility*. However, if we do not know anything about her overall preferences over the material payoffs, there is no utility function in the picture. Indeed, observe that it is perfectly conceivable to have different preferences than those just described above: for example, if Ann's preference relation is such that $5\$/5\$ > 6\$/4\$ > \cdots > 9\$/1\$ > 10\$/0\$ > 4\$/6\$ > \cdots > 0\$/10\$$, then we could infer that she cares for her well-being *and* also for Bob's well-being and her utility function should incorporate this. Hence, to conclude, this one observation does not give us enough information to meaningfully talk about Ann's utility function.

Considering that the argument in the previous paragraph can be extended (in an even more natural way) to games with a more complex structure than the Dictator Game, what could we say concerning our original question?

Bottom line: if we want to be able to sensibly state that a DM is rational, shouldn't we always divide an experiment from which we want to eventually draw conclusions on the rationality of the subjects in two stages? In stage 1 subjects' preferences should be elicited (in an incentive-compatible way) via their choices on all the possible subsets of the possible outcomes of the game that they have to play in stage 2, and only in stage 2 their actual choices in the game should be recorded and evaluated (along with their rationality... with the caveat that maybe it could not be possible to evaluate it, since stage 1 could show that some subjects' preferences simply cannot be represented via a utility function).

Piefrancesco Guarino
Maastricht University

A type of simulation which some experimental evidence suggests we don't live in

Do we live in a computer simulation? I will present an argument that the results of a certain experiment constitute empirical evidence that we do not live in, at least, one type of simulation. The type of simulation ruled out is very specific. Perhaps that is the price one must pay to make any kind of Popperian progress.

In electronics, a *soft error* is a type of error caused by a particle hitting a computer's memory banks. Early computer chips were manufactured with materials that emitted alpha particles due to radioactive decay. These alpha particles could hit memory cells and change memory values. The same phenomenon can happen if a cosmic ray hits the computer.

Suppose we live in a computer simulation with the following *x-$\hat{x}$ property*: for each memory-bit x in any computer in our world, there is a memory-bit $\hat{x}$ in the simulating computer such that $\hat{x}$ is used to store which value is stored in x. Then any such x is subject to two different types of soft errors:

- (Internal) Soft errors caused by simulated particles hitting x in our simulated universe.
- (External) Soft errors caused by real particles hitting $\hat{x}$ in the universe where the simulation takes place.

Further, assume the following *uni-directional property*: putting a simulated memory-bit inside a simulated vault does not protect it from external soft errors. We mean "vault" literally: a non-metaphorical barrier of hard matter in the simulated universe.

Putting a simulated memory-bit in a simulated vault might protect it from internal soft errors, because a thick vault might physically block incoming particles. The uni-directional property says this defense cannot prevent external soft errors. If we live in an external-soft-error-prone simulation with the x-$\hat{x}$ property and the uni-directional property, no vault we build can perfectly protect a memory-bank from all soft errors, because each memory-bit x in that memory-bank remains susceptible to external soft errors caused by real particles hitting $\hat{x}$.

A paper by O'Gorman et al (1996: Field testing for cosmic ray soft errors in semiconductor memories, http://ieeexplore.ieee.org/document/5389436/, IBM Journal of Research and Development **40** (1), 41–50) describes (p. 46) the following experiment and its results. A total of 864 modules were first run on the second floor of a two-story building for 4,671 hours, during which time, 24 soft errors were detected. Then, the same 864 modules were run for 5,863 hours in a nearby vault shielded by about 20m of rock, during which time, zero soft errors were detected.

The above results suggest that a vault of 20m of rock blocked all soft errors. By the above remarks, this is experimental evidence that we do not live in an external-soft-error-prone simulation with the x-$\hat{x}$ property and the uni-directional property. If we do live in such a simulation, then it should not be possible to protect a simulated memory-bank with a simulated vault.

Of course, this is not a mathematical proof, merely empirical evidence. The evidence could be improved, or the thesis falsified, with further experiments. What if we repeat the experiment and soft errors are detected in the vault? Without additional technology, we are unable to tell which soft errors were simulated and which were real. (We can only distinguish them vacuously: if zero soft errors occur, then zero are simulated and zero are real.) If soft errors persist in settings more and more hostile to internal-soft-errors, that is evidence that either we're overlooking (and failing to control for) some unknown source of internal soft errors, or else that external soft errors exist. The latter would entail we live in a external-soft-error-prone simulation, albeit not necessarily one with the x-$\hat{x}$ and uni-directional properties (maybe x-$\hat{x}$ fails for other memory-bits besides the ones tested; maybe soft errors in the simulating computer affect non-memory components of the simulated computers, indirectly manifesting as soft errors in simulated computers; and so on).

Perhaps this paper's most interesting conclusion is just that a non-contrived simulation hypothesis is falsifiable in a concrete way. One can easily imagine many types of simulations we could live in that are not external-soft-error-prone, or that lack the x-$\hat{x}$ property, or that lack the uni-directional property. I hope my argument will inspire falsifiable predictions of other types of simulations.

Samuel Alexander

News

Explanation and Understanding, Ghent, 23–25 May 2018

The workshop *Explanation and Understanding* was held on May 23rd to 25th at the Royal Academy of Dutch Language and Literature in Ghent, Belgium. It was the seventh in the *Logic, Reasoning and Rationality* series of workshops supported by the Research Foundation Flanders (FWO) through the scientific research network on Logical and Methodological Analysis of Scientific Reasoning Processes. The network brings together research groups from nine European universities carrying out research on relevant topics: Adam Mickiewicz University Poznań, Free University of Brussels, Ghent University, Ruhr-University Bochum, Tilburg University, University College London, University of Antwerp, Utrecht University and Vrije Universiteit Amsterdam. For the duration of the project, from 2016 till 2020, there are two workshops organized per year, one in spring and one in autumn.

The seventh workshop was organized by the Centre for Logic and Philosophy of Science of Ghent University, which coordinates the activities of the network, and the Department of Philosophy of the Vrije Universiteit Amsterdam, The Netherlands. The theme of the workshop, explanation and understanding, is one that is attracting a lot of attention in current philosophy of science. While most philosophers agree that science aims for explanation, and that one of the aims of explanation is to provide understanding, the unresolved question is how to account for this link between explanation and understanding. Most earlier accounts of explanation either took understanding to be of little philosophical interest—as merely a psychological phenomenon without epistemic import—or they considered the notion redundant—arguing that everything philosophically interesting could be captured with explanatory concepts, making ideas on understanding superfluous. Either way, for a long time the philosophy of science by and large ignored understanding, focusing on explanation instead. Things are very different now: the last two decades have witnessed a flourishing philosophical literature on the relation between explanation and understanding. And perspectives diverge greatly: while some argue that current ideas about understanding offer nothing distinctively new and relevant over and above explanatory concepts, others stress the distinctive role of understanding vis-á-vis explanation, and yet others make the case that one can have understanding without explanation.

Most of this recent literature appears to focus on full-fledged theories and explanations, neglecting other aspects of the explanatory enterprise. For instance, how does the idea of understanding play out in the context of explorative phases preceding scientific discovery, or in the early phases of discovery? What is the epistemic role of understanding in the usage of how-possibly models or hypothetical models to explain aspects of real systems? How does understanding relate to abstraction and idealization practices?

The workshop brought together 21 participants who presented talks on a great variety of topics related to the main theme. A number of talks discussed general issues, such as the relation between explanation and understanding, ontic versus epistemic conceptions of explanation, and the possibility of understanding without explanation. Other issues that were discussed included specific types of explanations (e.g., minimal structure explanations, narrative explanations) and the understanding they provide and the issue of whether, and if so how, various types of models—such as how-possibly models and machine learning models—can provide understanding. Finally, a number of talks discussed how understanding is achieved in different scientific disciplines and fields, such as evolutionary biology, population genetics, quantum mechanics, medical science, mathematics, history, ecology, and criminal law.

The keynotes (summarised below) were delivered by Matteo Colombo (Tilburg), Caterina Marchionni (Helsinki) and Alexander Reutlinger (Munich).

The backdrop of Matteo Colombo?s keynote, “I know that I know nothing. Explanation, Prejudice, and Intellectual Humility??, is the robust phenomenon that people are prejudiced towards members of groups with a worldview they perceive to be dissimilar from their own. In his talk, Matteo brought together ideas and methods from existing literatures on explanation and intellectual humility, and discussed how people’s intellectually humble explanatory reasoning might impact the relationship between dissimilarity and prejudice.

In her keynote, “Explanatory norms as frictions to integration: the case of economics and its neighbours??, Caterina Marchionni looked at the relationship between economics and neighbouring fields, and examined ways in which field-specific norms about explanation hinder the integration of mechanistic models across fields. She argued that the mechanism-based unity of science championed by Craver and other mechanistic philosophers is better captured by the image of a (dis-unified) cubist painting than that of a (unified) mosaic.

In the keynote, “Understanding and Non-Causal Explanation??, Alexander Reutlinger, argued that there is a unified account of causal and non-causal explanations, viz., the counterfactual theory of explanation, and he elaborated how this theory of explanation could provide a fruitful building block for a unified view of causal and non-causal modes of understanding.

Dingmar van Eck
Ghent
Erik Weber
Ghent
Henk de Regt
Vrije Universiteit Amsterdam

Calls for Papers

Formalization of Arguments: special issue of *Dialectica*, deadline 31 July.
Reliability: special issue of *Synthese*, deadline 11 November.
Instrumentalism about Epistemic Rationality: For and Against: special issue of *Synthese*, deadline 30 October.

What’s Hot in . . .

Medieval Reasoning

Anyone passingly familiar with Ancient or Medieval Philosophy has probably noticed that those thinkers had a fondness for making categorisations and distinctions – along with the same cutthroat polemical attitude and penchant for splitting hairs that most philosophers still have today. The traditional debates about the classifications and divisions of logic are numerous, extensive and (in)famous enough to ring a bell even with non-specialists: is logic a tool or a part of philosophy?; is it a scientific discipline (*scientia*) or a practical technique (*ars*)?; is it primarily about language (*scientia sermocinalis*) or about reason (*scientia rationalis*)?; what is its subject? The list could go on and on. It is easy to guess that the answers to these questions are often not clear-cut and tend to intermingle. But, overall, the consensus is that traditional logic is wider than contemporary logic and includes several ontological, linguistic, epistemic and even psychological issues that many of us wouldn’t even consider to be logic in the first place. In a sense, a lot of traditional logic deals with meta-logical questions, with problems pertaining to the philosophy of logic, and with what we would define as ”reasoning”. As far as the ”reasoning” aspect is concerned, it would be incorrect to believe that it was exhaustive of the traditional logical practices and theories, or that logicians of the past were unaware of any difference between ordinary reasoning, rationality, and what we would call logic proper as early as Aristotle. While an explicit distinction between *natural* and *artificial* logic already appears in the *Auctoritates Aristotelis*, a medieval anthology of theses more or less correctly attributed to Aristotle, the debates about the definition and the articulations of *logica naturalis* and *logica artificialis* became heated in the Later Middle Ages and continued to flourish throughout Early Modernity well into the 18th century – all the way up to Wolff and Kant. A comprehensive overview of the history of this distinction can be found in: M. Honen, "From Natural Thinking to Scientific Reasoning: Concepts of logica naturalis and logica artificialis in Late-Medieval and Early-Modern Thought", *Bulletin de philosophie médiévale* 52 (January 2010), p. 81-116. To make a long story short, from the different characterisations, distinctions, terminological uses, and feuds between opposing school of thought, a twofold picture of the use of the term "logic" emerges. On the one hand, we always find "artificial logic", which is the logic that is taught in school, put down in handbooks and treatises, and which needs to be studied and learned to be mastered; every author agrees that this is "logic proper". On the other hand, we find "natural" and "habitual" (*habitualis* or *usualis*) logic – respectively the principles of innate rationality and the ordinary reasoning procedures that are developed and refined through practice but do not require any formal learning setting. Overall, habitual logic comes close to be the late medieval analogue of what nowadays we call informal reasoning, even if many topics and inferential patterns that we would consider pertaining to informal reasoning were treated within the scope of artificial logic. However, the connection between habitual logic and logic in a proper sense was just as discussed as the connection between reasoning and logic is nowadays: while several philosophers believed that "logic" was used equivocally in the two cases, some remarked a closer analogical use, and others went for full univocity – holding that artificial logic is simply a systematisation and a refinement (or a formalisation) of habitual logic. While these debates are not as popular or widely studied as the ones mentioned above, they are still definitely worth a closer look.

Graziana Ciola
Philosophy, Scuola Normale Superiore, Pisa

Uncertain Reasoning

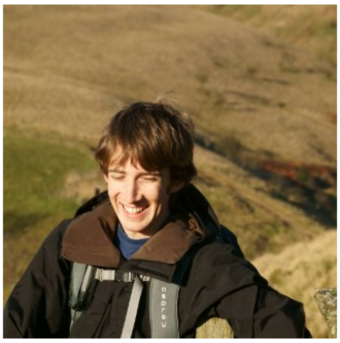

I'm sure we have all had that experience of seeing an actor in some film or TV show, and then suddenly they appear to be in everything you watch. A couple of years ago, every film out there starred Ryan Gosling. I recently had that experience not with an actor, but with a strange quirk of probability theory. I was preparing for a lecture on causation and causal inference in my philosophy of science class, and one of the things I read was Nancy Cartwright's "Causal Laws and Effective Strategies" (Noûs, 1979), in which Cartwright discusses "Simpson's paradox". I had heard of Simpson's paradox before, but only as a quirk of probability: it was interesting to see it put to use in a philosophical argument. The next week I was reading a paper about updating imprecise probabilities (Gong and Meng "Judicious Judgement Meets Unsettling Updating" arXiv:1712.08946v1) and again I was confronted with Simpson's paradox. A few weeks after that I was thinking about confirmation and I realised that again Simpson's paradox seemed relevant. Fitelson fleshes out this case ("Confirmation, Causation and Simpson's Paradox" Episteme 2017). None of these encouters with Simpson's paradox is particularly noteworthy in itself, but it does seem remarkable that it shows up in all these contexts; and to be confronted with them all within such a short space of time was surprising.

So what is Simpson's paradox? The classic example involves university admissions (the example comes from Cartwright). Berkley Graduate School (BGS) appeared to be admitting a higher proportion of its male applicants than its female applicants, which was a cause for concern for the school. This negative correlation between being female and BGS acceptance seemed like it might point to discrimination in admissions. However, for almost every department in BGS, there was a positive correlation or no correlation between being female and acceptance. That is, the direction of the correlation was reversed when you zoomed down to the level of the individual departments. This is the core of the Simpson's paradox: you can have a negative correlation between two variables, and yet if you partition the space, there is a positive correlation between the variables in each cell of the partition. What explains this weird behaviour is that the female students tended to apply to the departments with lower acceptance rates.

Cartwright uses Simpson's paradox to emphasise that we need to exploit our causal understanding of the situation to work out which correlations to take seriously in our theorising. Gong and Meng argue that Simpson's paradox suggests that care is needed when aggregating data. For example, if men's and women's applications had been randomly assigned to BGS departments, one would expect that a robust correlation at the department level would translate into a correlation in the same direction at the level of the whole school. Fitelson uses a diagnosis of why Simpson's paradox seems paradoxical to make a subtle point about two ways of understanding confirmation.

Simpson's paradox is a fun little quirk of probability theory that has been put to a remarkable variety of dialectical purposes: Simpson's paradox is the Ryan Gosling of probability theory.

Seamus Bradley
Philosophy, University of Tilburg

Mathematical Philosophy

Non-causal explanations are a hot topic in current mathematical philosophy, philosophy of science, and philosophy of mathematics.

Providing philosophical theories of scientific explanations has been a central task for philosophy of science, from its beginnings in the early 20th century until the present day. For the past three decades, *causal accounts* of scientific explanations have been the received view of explanation and the detailed philosophical analysis of causal explanations has been the main focus of the philosophy of scientific explanation. According to causal accounts, the sciences explain by identifying the causes of the phenomenon to be explained. In the present debate, indeed hardly anyone denies the significance and epistemic value of causal explanations in the sciences.

However, a significant change has been taking place in philosophy of science since the mid 2000s. The received causal view has been called into question. Several philosophers have raised the objection that causal accounts of explanation fail to provide a complete picture of explanations in the sciences. In particular, causal accounts do not do justice to a large number of, what at least seem to be, compelling examples of *non-causal explanations* – that is, scientific explanations whose explanatory power does *not* derive from identifying the causes of the explanandum phenomenon at issue (for surveys to the recent literature and for paradigmatic examples of non-causal explanations in the sciences, see Reutlinger, A. [2017] ("Explanation Beyond Causation? New Directions in the Philosophy of Scientific Explanation", *Philosophy Compass*, Online First) and Mancosu, P. [2018] ("Explanation in Mathematics", *The Stanford Encyclopedia of Philosophy*, Summer 2018 Edition, Edward N. Zalta (ed.), URL = https://plato.stanford.edu/archives/sum2018/entries/mathematics-explanation/). The existence of non-causal explanations creates a serious challenge to causal accounts: if there are non-causal explanations, then causal accounts can no longer be taken to specify necessary and sufficient conditions for explanation.

The discussion of non-causal explanations is also motivated by explanatory practices in 'extra-scientific' domains. After all, the natural and social sciences are not the only epistemic projects striving for explanation. Most importantly, proofs in pure mathematics are at least sometimes taken to be explanatory. However, if mathematical proofs possess explanatory power, then such proofs are taken to explain in a *non-causal* way why a particular theorem holds – see, for instance, Colyvan, M. [2012] (*An Introduction to the Philosophy of Mathematics*, Cambridge: Cambridge University Press), Lange, M. [2016] (*Because Without Cause. Non-Causal Explanations in Science and Mathematics*, New York: Oxford University Press), and Mancosu, P. [2018] ("Explanation in Mathematics", *The Stanford Encyclopedia of Philosophy*, Summer 2018 Edition, Edward N. Zalta (ed.), URL = https://plato.stanford.edu/archives/sum2018/entries/mathematics-explanation/).

Examples of non-causal explanations in the sciences and in mathematics raise challenging questions such as the following:

1. Supposing that there are two types of explanation in the sciences (causal and non-causal ones), is it possible to defend a monist approach to scientific explanation? That is, is it possible to develop one single philosophical theory of explanation capturing both causal and non-causal explanations? Or is one forced to accept explanatory pluralism (that is, the view that there is no such single theory and one needs to appeal to two or more accounts of explanation to capture causal and non-causal explanations)?

2. Is it possible to extend accounts of non-causal explanation in the sciences to non-causal explanations in pure mathematics?

3. What are convincing criteria for distinguishing between causal and non-causal explanations in science and mathematics?

4. Can accounts of non-causal explanations overcome the problems troubling Carl Hempel's covering-law model, given that the standard solutions to these problems seem to depend on appealing to causation? For instance, how does one account for the explanatory asymmetry in the case of non-causal explanations (supposing that non-causal explanations generally display such an asymmetry)?

Addressing these questions is a crucial task for future research on non-causal explanations in science and mathematics. Lange, M. [2016] (*Because Without Cause. Non-Causal Explanations in Science and Mathematics*, New York: Oxford University Press) and the essays in Reutlinger, A. and Saatsi, J. (eds.) [forthcoming] (*Explanation Beyond Causation. Philosophical Perspectives on Non-Causal Explanations*, Oxford: Oxford University Press) constitute current attempts to make progress on these issues. The topic of non-causal explanations is not only worth dealing with, it is also a great opportunity for mathematical philosophers, philosophers of science and philosophers of mathematics to collaborate more closely.

Alexander Reutlinger
Munich Centre for Mathematical Philosophy

Evidence-Based Medicine

Should we eat more fat, or carbohydrates, or maybe even more protein? Macronutrient composition is always a hot topic, and the evidence base for dietary guidelines is constantly in flux. This is exemplified by the recent retraction of a clinical trial that originally seemed to show the dramatic positive effects of a high-er fat diet on cardiovascular health. The PREDIMED trial concluded that following a 'Mediterranean diet' reduced the chance of mortality through cardiovascular disease by up to 30%. The Mediterranean diet consists of a high intake of unsaturated fats mainly from olive oils, fish and nuts, a type of fat that is linked with positive effects on cardiovascular health. The results of the study are now in doubt, due to the way in which the study was carried out. Of particular importance was the finding that 21% of the participants were not properly randomized. In a randomized control trial, this does make one start to doubt the veracity of the results. The paper has been re-analysed and re-published. After making some statistical adjustments, weaker yet still positive conclusions have been made. There is dispute over whether the implementation problems of the study negate the conclusions entirely, with experts disagreeing over the matter.

This raises some important methodological questions. Can statistical adjustments paper over the faults of implementation? Why bother to implement a study well if we can just do some statistical adjustments once we realise the study was flawed? Can other evidence support conclusions from flawed studies? The authors of the study support their conclusions by referring to the wealth of evidence supporting such a diet from systematic reviews of other studies. However, as this PLOS blog post points out, they have been very selective in which systematic reviews they have chosen as support. Systematic reviews can be as unreliable as other study designs. Many high-quality reviews were not included in the PREDIMED list including ones done by the Cochrane Collaboration and the UK's National Institute for Clinical Excellence (NICE). Importantly, "2 Cochrane reviews rejected the PREDIMED trial from consideration, even before these revelations, primarily because of taking issue with the control group." Reviewers at NICE also assessed the PREDIMED trial as at serious risk of bias, and providing low or very low quality data. Even with these reviewer's concerns, the trial had gained prominence as 'support' for the mediterranean diet, showing the importance of reporting on results of systematic reviews as well as of individual trials, and always keeping quality in mind.

The PREDIMED trial is specifically referred to in a recent article in the BMJ summarising the state of the evidence regarding fat and health. This article went to press before the retraction, but still refers to doubt surrounding the PREDIMED trial. From the litany of methodological problems mentioned in this article that dietary science research faces, it is no wonder that the 'conclusive' results of the PREDIMED trial came to be doubted. Good evidence does seem to suggest that Low Density Lipoprotein (LDL) is a cause of coronary heart disease, and that high levels of saturated fat intake in ones diet increases levels of LDL in the blood. Putting the two together allows the inference that high saturated fat intake promotes coronary heart disease. Strangely, high-quality evidence showing that decreasing saturated fat intake in one's diet will lower the risk of heart disease is lacking. However, health guidelines do promote replacement of saturated fats with unsaturated fats. This is done on the basis of a synthesis of lower-quality kinds of evidence, like observational studies, with large RCT's providing only moderate-quality evidence (due to methodological flaws), and good evidence from laboratory science showing the link between LDL and coronary heart disease. This may go against typical EBM guidelines, but is necessary to do for such an important topic, with many significant health, economic and social consequences.

Instead of such a synthesis, the usual advice is to continue to try to obtain high-quality evidence from higher quality study designs (like RCTs). This approach may not be of use in this situation. The PREDIMED trial is only one instance of a large scale RCT being beset by significant quality problems. In general, dietary science research needs to run trials on 1000's if not 10,000's of participants over long time scales, as the negative health effects of certain diets take a long time to manifest. One major problem facing such trials is something as simple as dietary adherence - ensuring the participants keep to one diet over a span of up to 30 years seems to be an insurmountable difficulty. The potential for confounders to bias the results is also significant over such large time frames. Most of the long-term trials we have were conducted over a period of time in the late 20th Century when public health improved for many external and separate reasons. A modest replacement of saturated fats with unsaturated fats in this time period was associated with a 75% reduction in coronary heart disease. It was also associated with dramatic drops in smoking rates, an increasing focus on personal fitness, and major improvements in medical science. This is something that is compounded by the complexity of nutrition. One potential issue is the way in which nutrients interact: saturated fat in one food-stuff may have different effects on health than in other food-stuffs due to the presence of other nutrients. The list of issues is not limited to these problems, and the interested reader is directed to the BMJ article.

Methodological problems when studying such a complex issue seem unavoidable. The silver-bullet of randomization and double-masking appears to be ineffective against pervasive confounding. Evidence synthesis of lower-quality forms of evidence to inform guidelines is controversial, but such guidelines rely on these forms of evidence to make the inferences that they do. Better acknowledgement of the vital role this evidence plays may improve dietary guidelines by avoiding over-reliance on trials such as PREDIMED, which may always be flawed.

D.J. Auker-Howlett
Philosophy, Kent

Events

July

PoLaL: Workshop in the Philosophy of Language and Logic, Humboldt University of Berlin, 2–3 July.
EpiPet: Epistemic Peterification Conference, University of Bristol, 2–3 July.
AAoL: Australasian Association of Logic Meeting, Victoria University of Wellington, 6–7 July.
PaR: Proofs and Representations, Munich, 6–8 July.
VEW: Vice Epistemology Workshop, University of Sheffield, 9 July.
Women in the History of Science, Philosophy and Literature, Syros, Greece: 12–13 July,

.

ESoN-EF: The Epistemic Significance of Non-Epistemic Factors, University Osnabrück, Germany, 12–14 July.
Dcsn&AI: Decision Theory & the Future of Artificial Intelligence, Munich, 27–28 July.

August

VoMA: Varieties of Mathematical Abstraction, University of Vienna, 1–3 August.
TaE: Workshop on Time and Explanation, Milan, 20–21 August.

September

PLP: The 5th Workshop on Probabilistic Logic Programming, Ferrara, Italy, 1 September.
WAW: Warsaw Argumentation Week, Warsaw, Poland, 6–16 September.
EPINON: Epistemology in Ontologies, Cape Town, South Africa, 17–18 September.

Courses and Programmes

Courses

LUCG: Logic, uncertainty and games, Como, 9–13 July.
SIPTA: 8th School on Imprecise Probabilities, Oviedo, 24–28 July.

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 Epistemology, California State University, Fullerton, deadline open.
Lecturer: in Epistemology/Metaphysics, University of Lincoln, deadline 15 July.
Fellow: in Philosophy of Science, University of Bern, deadline 16 July.
Lecturer: in Logic and Decision Theory, University of York, deadline 20 July.
Assistant Professor: in Kant; Philosophy of Science, University of Notre Dame, deadline 26 July.
Associate Professor: in Philosophy of Science, University of Oslo, deadline 10 August.

Studentships

PhD positions: in Logical Methods in Computer Science, Austria, deadline 1 July.
PhD position: in Logic, University of Bonn, deadline 31 July.