British Society for the Philosophy of Science

Annual Conference
Manchester, 2-3 July 2015

Contents:

Programme	3
About the 'Meet the editors' session	4
Open sessions timetable	5-6
Information for delegates:	7-8
Arrival, checking in and checking out	
Wifi access	
Places to eat and drink	
Things to see and do	
Information about the open sessions:	9
Timing and organisation	
Computers and data projectors	
Events policy	
Open sessions abstracts	10-98
Local map	99

PROGRAMME

Thursday 2 July

10.00-11.00	Registration/tea & coffee
11.00-12.30	Plenary I: Mauricio Suarez, ‘Propensities and statistical modeling’ (chaired by Hugh Mellor)
12.30-1.30	Lunch
1.30-3.00	Open Session I
3.00-3.30	Tea & coffee
3.30-5.00	Open Session II
5.00-5.30	Tea & coffee
5.30-7.00	Plenary II: Havi Carel, ‘“If I had to live like you, I think I'd kill myself": illness, wellbeing and medical practice’ (chaired by Peter Clark)
7.00-8.00	Drinks reception, kindly sponsored by OUP
8.00-9.30	Conference dinner (Mumford Room, Manchester Meeting Place)

Friday 3 July

8.30-9.00	Tea & coffee
9.00-10.30	Open Session III
10.30-11.00	Tea & coffee
11.00-12.30	Plenary III: Kim Sterelny, ‘Cumulative cultural evolution and the origins of language’ (chaired by Steven French)
12.30-1.00	BSPS AGM: Harwood Room. All welcome. Come and have your say on how your society is run!
1.00-2.15	Lunch
1.45-2.15	BJPS ‘meet the editors’ session: Harwood Room
2.15-3.15	Open Session IV
3.30-4.30	Open Session V
4.30-5.00	Tea & coffee
5.00-6.30	Plenary IV: Katherine Brading, ‘Émilie du Châtelet and the foundations of physical science’ (chaired by Harvey Brown)

Venues

Plenaries: Harwood Room, Barnes Wallis Building

Refreshments/lunch/drinks reception: Barnes Wallis Room, BWB

Open sessions: Harwood Room; rooms 1, 2, 3 and 8, Manchester Meeting Place

About the ‘Meet the editors’ session (Friday 1.45)

Come along and meet the editors of the British Journal for the Philosophy of Science. Find out how the journal works, what the various stages are that a submission goes through and how to respond to referees' reports (including 'cursing their ancestors and descendants').

This will be primarily a 'Q&A' session and we're particularly keen to respond to questions and concerns from postgrads, postdocs and early career folk.

OPEN SESSIONS TIMETABLE

Harwood	Room 1	Room 2	Room 3	Room 8
Open Session I: Thursday, 1.30pm – 3.00pm				
Chair: Harvey Brown	Chair: Bryan Roberts	Chair: Liz Irvine	Chair: Dana Tulodziecki	Chair: Hugh Mellor
Olivier Sartenaer. Emergent quasiparticles: the case of the fractional quantum Hall effect	Alexander Franklin. Universality Explained?	Alexander Gebharter. Causal exclusion and causal Bayes nets	Kevin Coffey. Reconsidering unconceived alternatives: prospects for scientific realism	Brice Bantegnie. A shift in focus: from mental states to mental capacities.
Matthias Egg. Do we need a primitive ontology to make quantum mechanics empirically coherent?	Vincent Ardourel and Julie Jebeile. Are numerical solutions preferable to analytical solutions?	Michael Baumgartner and Lorenzo Casini. Establishing constitutional relations, in theory and in practice	Peter Vickers. No miracles? Scientific realism and the 1811 gill slit prediction	Rosa Hardt. The interdependence of emotion and sensory experience
Marton Gomori, Laszlo E. Szabo and Zalan Gyenis. Operationalist approach to quantum theory: two representation theorems	Michael Miller. Exact models and physical semantics	Tobias Starzak. Morgan's Canon: interpretation and justification	Robert Northcott. Approximate truth and scientific realism	Daniel Calder. Ramsey reconsidered: applying the job-descripton challenge to contemporary cognitive science.
Open Session II: Thursday, 3.30pm – 5.00pm				
Chair: Steven French	Chair: Kerry McKenzie	Chair: Richard Pettigrew	Chair: Flavia Padovani	Chair: Elselijn Kingma
Bryan W. Roberts. Geometrizing quantum theory	Erik Curiel. If metrical structure were not dynamical, counterfactuals in general relativity would be easy	Juha Saatsi and Lina Jansson. Varieties of abstract explanations: causal, non-causal, and mathematical	Dana Tulodziecki. From zymes to germs: discarding the realist/anti-realist framework	Janette Dinishak. Autism, aspect-perception, and deficit explanations of human differences
Raymond Lal. The topology of contextuality: a unifying concept in quantum theory and logic	Joshua Eisenthal. The problem of space	Harjit Bhogal. Three dimensions of explantory goodness	Haixin Dang. Theory choice during conceptual change: The case of WH Bragg and X-rays	Rachel Cooper. The unluckiness of the disordered
Davide Romano. The meaning of the mass in Bohm's theory and classical mechanics: a case study from the classical limit	Dennis Lehmkuhl. The neighborhood of general relativity in the space of (spacetime?) theories	Matteo Colombo and Jan Sprenger. Explanatory value and probabilistic reasoning: an empirical study	David Schroeren. Theoretical equivalence as explanatory equivalence	Magdalena Antrobus. Good grief: epistemic and psychological benefits of depressive mood

Harwood	Room 1	Room 2	Room 3	Room 8
Open Session III: Friday, 9.00am – 10.30am				
Chair: Oliver Pooley	Chair: Bryan Roberts	Chair: Liz Irvine	Chair: Peter Clark	Chair: Rachel Cooper
Tomasz Placek. Indeterminism and bifurcating geodesics	Thomas Moller-Nielsen. Weak discernibility, again	Hugh Desmond. Natural selection: convergence and causality	Casey Helgeson. Low confidence in extreme probabilities	Ana-Maria Cretu. What good is realism about natural kinds?
Radin Dardashti. No alternatives for what? Non-empirical evidence in the case of string theory	Samuel Fletcher. Limits of Nagelian reduction	Zachary Ardern. Evolution and causal role functions	Benjamin Bewersdorf. Conceptual Learning and Bayesian epistemology	Joe Dewhurst. Natural kinds and folk kinds in the psychological sciences
Lena Zuchowski. A sideways glance at Smale's fourteenth problem: definition and ontology of chaos	Alastair Wilson. Naturalizing recombination	Brian Garvey. The evolution of morality and its rollback	Jürgen Landes and Jon Williamson. How an objective Bayesian integrates data	Kerry McKenzie. Intrinsicality and the Goldilocks Principle: fundamentality as an untapped resource for structuralism
Open Session IV: Friday, 2.15pm – 3.15pm				
Chair: Oliver Pooley	Chair: Phyllis Illari	Chair: Janette Dinishak	Chair: Lena Zuchowski	Chair: Bryan Roberts
Flavia Padovani. Coordination, measurement, and the problem of representation of physical quantities	Elselijn Kingma. Functions at the interface of biology and technology: synthetic biology, health and disease	Anna-Maria Asunta Eder. In defense of a credence interpretation of probability	Nicolas Wüthrich. Conceptualizing uncertainty: An assessment of the latest uncertainty framework of the Intergovernmental Panel on Climate Change	Seamus Bradley, Karim Thébault and Alexander Reutlinger. Modelling inequality
Neil Dewar. Symmetry, differences, and naturalism	Karen Kovaka. Rejecting replicators	Nick Tosh. Ensemble realism: a new approach to statistical mechanical probability	Marina Baldissera. In what sense is uncertainty intrinsic to climate science?	James Fraser. Groundwork for a neo-Galilean approach to idealisation
Open Session V: Friday, 3.30pm – 4.30pm				
Chair: Steven French	Chair: Peter Clark	Chair: Phyllis Illari	Chair: Richard Pettigrew	Chair: Lena Zuchowski
Owen Maroney and Daniel Bedingham. A flash, a collapse and a boundary condition: where did the asymmetry come from?	Callum Duguid. Best systems accounts and metalaws	Christopher Blunt. How to create false positives and influence people: cohort multiple RCTs and the grades of recommendation	Marie Barnett. Reasons and conditional preferences	Veli-Pekka Parkkinen. Mechanism-based extrapolation reconsidered
Chris Timpson. Bell's theorem, local causality, explanation and Everett		Ioannis Votsis. How to make a long theory short: lessons from confirmation	Gregory Wheeler and Conor Mayo-Wilson. Epistemic decision theory's reckoning	Viorel Paslaru. Integrative pluralism in scientific explanations, and a lesson from ecology

Information for delegates

Arrival, checking in and checking out

The map at the end of this information pack shows a route from Piccadilly Station to the conference venue (Barnes Wallis Building/Manchester Meeting Place). The easiest way to get straight to the conference accommodation at Weston Hall from Piccadilly is to walk down Granby Row until you hit Sackville Street. Then turn left, go under the railway bridge and cross over Charles Street; the Weston Hall entrance will then be set back from the road on your right after you've passed the first block.

Check-in at Weston Hall is after 2pm. If you're arriving for the start of the conference on Thursday, we can store your luggage in a locked room right by the Barnes Wallis Room; just ask at the conference reception desk. Please retrieve any luggage by the beginning of the drinks reception at 7pm at the latest.

Reception at Weston Hall is open 24 hours.

You will need to check out of your room by 10am on your day of departure. Again, we can store your luggage; just ask at reception.

Wifi access

Wifi access via eduroam is available in the communal areas of Weston Hall. Access via eduroam or a conference guest account is available at the conference venue.

To access eduroam, your username is your own username that you use to log in to your university's IT services, followed by '@xxx.ac.uk' (or '@xxx.edu, or whatever) and your normal email password. NB this means that your normal university email address itself normally won't work as your username. (E.g Freya Bloggs at Manchester would have something like 'msgssfb1@manchester.ac.uk' and not 'f.bloggs@manchester.ac.uk'.)

If you can't access eduroam for any reason, please ask for a guest account at the conference reception desk.

Places to eat and drink

Nearby cafes: There is a café area in the Barnes Wallis Building, open 10am to 2.30pm daily.

There's also Starbucks on Sackville Street, very close to Weston Hall; and Olive on the corner of Whitworth and Sackville. If you need really good coffee, you're probably going to have to trek to North Tea Power in the Northern Quarter (36 Tib Street, on the first floor).

Nearby pubs: Lass O'Gowrie (or – probably not as good – Joshua Brooks), Charles St, just round the corner from Weston Hall; The Bull's Head, corner of London Road and Granby Row (opposite Piccadilly Station).

A little further away in the city centre, you might try Sam's Chop House, Mr Thomas's Chop House, or Sinclair's Oyster Bar – all classic old-style pubs. If you prefer bars, there are zillions on Canal Street – running along the north bank of the canal for about three blocks north-eastwards starting at Princess Street. A nice place to sit outside by the canal if the weather's good.

Restaurants: If you're looking for a good cheap alternative to the conference dinner, there's the Curry Mile on Oxford Road in Rusholme. (Catch nearly any bus heading south down Oxford Road – check with the driver that it goes through Rusholme. You'll know when you've got there.) Or there's Chinatown in the city centre: the few blocks just to the north-east of Portland Street and Princess Street. Or – very cheap but great – there's Habesha, an Ethiopian restaurant on the corner of Sackville and Richmond (upstairs above Istanbul Express).

Things to see and do

Art galleries: Manchester Art Gallery in the city centre, or – if you don't mind a short bus ride – the newly extended and splendid Whitworth Art Gallery. (Catch nearly any bus heading south down Oxford Road – check with the driver that it goes through Rusholme. It's a big Victorian building on your right, on the edge of Whitworth park, just north of Rusholme. If you get to the Curry Mile, you've gone too far.) There's also the Lowry at Salford Quays – a tram ride away – if you like Lowry.

Museums: The Museum of Science and Industry, near Deansgate Locks, is a bit of a walk but well worth a visit. It has loads of kit, including a working replica of Baby, the world's first computer to store and run a programme (which was built at the University of Manchester) and lots of working steam engines. The museum is partly housed in the Manchester terminus of the world's first passenger railway, the Liverpool & Manchester. If you're very lucky, you might even get a (very short) ride on steam train. And you can walk through a Victorian sewer. The People's History Museum is good too.

All of the above have free admission.

Other: HOME Manchester is an arthouse cinema, gallery and café. On Whitworth Street, just west of Oxford Road on your left. Manchester Library and (next door to it) the Town Hall are worth looking at (and the library has exhibition areas as well), as is the John Rylands Library on Deansgate, which also has exhibition areas.

Finally, don't forget to go and say hello to Alan Turing, who's sitting on a bench in Sackville Gardens, right by the conference venue.

Information about the open sessions

Timing and organisation

Each slot is 30 minutes long. **The chair** should ensure that it starts promptly on the hour or half hour, and (apart from the last slot in each session) finishes a couple of minutes before the hour/half hour, to allow audiences to move between sessions.

Speakers should ensure that they finish their talk within 20 minutes, and may be cut off by the chair if they fail to do so in order to allow sufficient time for questions.

Speakers are politely requested to attend the whole session of which their talk is a part.

Audience members are requested to keep their questions very brief, only to ask one question at a time, and not to come back to the speaker with a follow-up question or remark unless explicitly permitted to do so by the chair. The 'hand/finger distinction' will not be deployed.

Computers and data projectors

These are available in all of the open session rooms – or you can plug the VGA adaptor into your own laptop (Mac users: please bring your own adaptor). If you are bringing your presentation on a memory stick, please arrive at the session a few minutes early and upload it onto the PC/laptop provided in advance.

Chairs are requested to arrive at their session a few minutes early to ensure that any presentations have been uploaded, that the data projector is on and working, and that all the speakers are present.

Events policy

All conferences organised by Philosophy at the University of Manchester operate an events policy, covering seminar conduct, harassment, and use of sexually explicit language and imagery. The policy can be found at:

www.socialsciences.manchester.ac.uk/subjects/philosophy/events-and-seminars/events-policy

Open Sessions abstracts (in alphabetical order)

Magdalena Antrobus. Good grief: epistemic and psychological benefits of depressive mood

In this paper, I ask whether depressive mood has epistemic and psychological benefits for the subject, and what is a character of the relation between such benefits.

Symptoms of depression are commonly perceived as psychological difficulties that compromise wellbeing. Considerations about the psychological costs of depressive mood and its adverse effects on functioning seem to rule out the possibility that 'feeling low' is linked to any benefits. Yet, it has been argued that in certain circumstances, depressive symptoms enhance more accurate beliefs about self and reality. One suggestion is that depressive mood enhances accuracy in respect of self-related judgments (depressive realism). This claim has been evidenced in numerous empirical studies over one's sense of control in contingency tasks (that is, the tasks in which the outcome might be perceived as a result of the subject's actions) (for ex. Alloy and Abramson, 1979; Abramson, Alloy and Rosoff, 1981; Dobson and Pusch, 1995; Presson and Benassi, 2003; Msetfi, Murphy, Simpson and Kornbrot, 2005). Arguably, 'feeling low' can enhance the accuracy of one's beliefs and help perceive the reality more precisely.

In this paper I endorse this claim in a light of the available evidence and argue, that by making a contribution to the acquisition of more accurate beliefs, depressive mood can be seen as epistemically beneficial.

Furthermore I argue that by contributing to better accuracy of beliefs with regards to self, depressive mood has also potential psychological benefits for the subject. Empirical evidence and real life observations show that accurate judgments of own capabilities can increase efficiency of self-defensive psychological strategy in a situation of experienced anxiety (defensive pessimism).

Defensive pessimism is a type of cognitive strategy performed by an agent in an attempt to take cognitive control over experienced anxiety, so that own performance is unimpaired (Norem and Cantor, 1986)

The notion of defence applies when the cognitive act of pessimism takes place in order to prevent greater harm from occurring (e.g. the level of anxiety becomes debilitating). The notion of pessimism applies to the cognitive strategy of setting unrealistically low expectations for the success of one's own performance (a situation that is a source of anxiety). Setting low expectations, in turn, plays a motivational role for the subject, who then takes an adequate action in order to protect herself from potentially greater psychological harm.

Here is an example of a defensive pessimism:

"Think, for instance, of straight-A students who have never failed a test in their lives but repeatedly insist that they are, without question, going to 'bomb' an upcoming exam. Nothing their friends can say reassures them; indeed, reminding them of their past success seems only to lead to more anxiety or confusion. These persons proceed to rush home, drink gallons of coffee, study furiously throughout the night and, annoyingly but not surprisingly, receive the highest score in the class. This success does not come without considerable effort devoted to preparation, however, and the anxiety, although perhaps unjustified, is very real." (Norem and Cantor, 1986b, p.1209).

Defensive pessimism has been evidenced as an effective cognitive strategy used in situations of high anxiety by healthy individuals, but largely ineffective when used by people suffering from clinical depression (Norem and Cantor, 1986a, 1986b). In the latter group, experiencing high levels of anxiety lead to increased depressive symptoms rather than motivates people to take a positive action. Here I argue that whilst severe forms of depression may prevent individuals from successful implementation of defensive pessimism strategy and whilst healthy individuals are often subjects to positive illusions, which might affect the efficiency of their actions, milder symptoms of depression (in form of depressive mood) can contribute to defensive pessimism success in such a way that thanks to more accurate perception of herself, the subject is likely to take a more effective action.

Re-considering the phenomenon of depressive mood in terms of its potential epistemic and psychological benefits leads to a more balanced view of the role of depressive symptoms in a person's cognitive and emotional life.

References

Abramson, L. Y., Alloy, L. B., & Rosoff, R. (1981). Depression and the generation of complex hypotheses in the judgement of contingency. Behaviour Research and Therapy, 19, 35–45. doi:10.1016/0005-7967(81)90110-8

Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: sadder but wiser? Journal of Experimental Psychology. General, 108(4), 441–485. doi:10.1037/0096-3445.108.4.441

Dobson, K. S., & Pusch, D. (1995). A test of the depressive realism hypothesis in clinically depressed subjects. Cognitive Therapy and Research, 19(2), 179–194. doi:10.1007/BF02229693

Msetfi, R. M., Murphy, R. A., Simpson, J., & Kornbrot, D. E. (2005). Depressive realism and outcome density bias in contingency judgments: the effect of the context and intertrial interval. Journal of Experimental Psychology. General, 134, 10–22. doi:10.1037/0096-3445.134.1.10

Norem, J. K., & Cantor, N. (1986a). Anticipatory and post hoc cushioning strategies: Optimism and defensive pessimism in "risky" situations. Cognitive Therapy and Research, 10(3), 347–362. doi:10.1007/BF01173471

Norem, J. K., & Cantor, N. (1986b). Defensive pessimism: harnessing anxiety as motivation. Journal of Personality and Social Psychology, 51(6), 1208–1217. doi:10.1037/0022-3514.51.6.1208

Presson, P. K., & Benassi, V. A. (2003). ARE DEPRESSIVE SYMPTOMS POSITIVELY OR NEGATIVELY ASSOCIATED WITH THE ILLUSION OF CONTROL? Social Behavior and Personality: An International Journal. doi:10.2224/sbp.2003.31.5.483

Zachary Ardern. Evolution and causal role functions

The standard view of biological functions amongst the evolutionary biology community appears to be that they are explicable solely in terms of an etiology of selected effects – a feature is 'functional' if it was retained in a population due to natural selection. Controversy over results from the ENCODE project which purported to show that most of the human genome is functional has highlighted the longstanding debate in biology over how, after Darwin, to make sense of the apparently teleological notion of function. I argue that results from evolutionary and molecular biology provide support for a 'causal role' view of biological functions, where the function of a feature is determined by the role played within the biological system.

The concept of a biological feature's function is often connected to the origin of that feature, but the questions of causal history and functional role are distinct and should not be conflated, as argued by Robert Cummins 40 years ago. One example of the distinction is given by the fact that natural selection often selects mutations with a 'loss of function' effect – for instance in the famous case of the mutations in the haemoglobin gene responsible for both sickle cell anaemia and resistance to malaria. It has been argued that such loss of function mutations are the primary method of adaptation observed in microbial experimental evolution. It is not clear that every result of natural selection should be considered 'functional' – as another example, in the extreme case natural selection leads to extinction. Further, on a causal role view, where a biological feature is functional if it contributes to the overall functioning of the organism, it is plausible to say that natural selection currently acts on the feature (i.e. acts to prevent its loss from the population) or will in future act on it – but it cannot thereby be concluded that the feature arose due to natural selection.

Conceptualising functions solely in terms of selected effects fails to make sense of many standard uses of 'function' in the biological literature, including key evolutionary concepts. Four examples: firstly, discussion of convergent evolution – where the same feature arises independently in different evolutionary lineages – assumes that it is coherent to talk about the same function having different selective histories, so function cannot be identified with selective history. Secondly, exaptation – where a biological structure undergoes a shift in function over the course of evolution – raises questions over what will count as the 'real' function of a structure. The causes of a structure being either maintained in a population or undergoing 'exaptation' into a current function are likely to be different to the causes of its origin. Thirdly, on a selected effects view, at least without further nuance regarding current functional role, it seems that vestigial organs and parasitic genetic elements will count as functional, regardless of their current contribution to the organism. Finally, it seems reasonable to talk about the potential functions of biological structures in novel environments and of molecular features in novel genomic contexts, presenting a challenge to purely backwards-looking views of function. The biological literature accepts that the same biological function can be achieved by different structures, and also that similar structures can arise via very different evolutionary histories.

It is generally thought that teleological language previously taken literally in terms of divine design can be retained due to the designer-substitute of natural selection. Darwinian etiological accounts of function are appealing because of the apparent legitimacy of 'goal associated' language. As we begin to understand natural selection and evolutionary processes in more detail, the appropriate domain of teleological language needs to be re-evaluated. The process of constructive neutral evolution has been posited to explain much cellular complexity, raising questions for functionality. Could a feature evolved through neutral processes such as this be functional? Additionally, if selection acts at levels other than the individual organism (as is almost certainly the case), are the products of such a process functional on a selected-effects view?

Finally, I note that inferring functions through experiment is possible if functions are causal roles, but often not possible if they can only be inferred in light of a structure's etiology. For instance, discerning whether in fact a particular mutation was fixed through neutral or selective processes is at best difficult, and in some cases impossible. Evolutionary and experimental considerations favour moving beyond a backwards-looking account of biological function.

Vincent Ardourel and Julie Jebeile. Are Numerical Solutions Preferable to Analytical Solutions?

An important task in empirical sciences like physics consists in getting numbers about the properties of the system under study, as well as its past, present and future physical states, via the resolution of differential equations. The resolutions are done either analytically or numerically. Either one uses an analytical method to express, when possible, an exact function of variables or one uses a numerical method to obtain approximate numerical results. Commonly, for the purpose of getting numbers, the analytical method is considered as the most valuable choice. As emphasized by Humphreys (2004, p. 64), one of the main reason is that analytical solutions are exact and this guarantees that the numbers obtained from them are consistent with the differential equations; whereas numerical methods often generate numerical errors on their own and therefore produce numbers which deviate — a little or too much — from the equations.

However, in some cases, scientists use numerical methods even though analytical solutions are available. For example the Schrödinger equation is generally solved numerically on a digital machine, whereas its analytical solutions are known (French and Taylor 1998, p. 174). This is a quite surprising practice since numerical solutions are generally viewed as less valuable than analytical solutions. In this paper, we tackle the question why scientists can prefer approximate numerical solutions over analytical solutions. This question has not received much philosophical attention until very recently (Fillion and Bangu 2015) while it may lead to important considerations about applied mathematics.

In order to answer the question, we exclude the obvious situation in which scientists use numerical methods because analytical solutions are not available. This situation includes the case where the equations are intractable per se — i.e. there are impossibility theorems that prove the mathematical absence of tractable solutions — and the case where scientists ignore whether the equations are analytically solvable. We rather focus on the most interesting cases where analytical solutions are available, but scientists still prefer to use numerical methods. Via several examples, we argue that their preference can be explained by the fact that, while getting numbers seems at first glance to be facilitated with analytical solutions, this task is actually often much easier with numerical methods. In other words, ease to get numbers sometimes prevails over exactness of solutions.

The paper is organized as follows. To begin with, we emphasize that analytical solutions are exact and make clear why, at first glance, they are more valuable than numerical solutions in empirical sciences. We then suggest three main reasons why numerical solutions are nevertheless sometimes preferable to analytical one for the purpose of getting numbers:

First, we argue that analytical solutions can make numerical applications difficult or impossible. Some properties of analytical solutions are in conflict with the requirement of making reliable numerical predictions. As an example, we focus on the case of the N-body problem. Although there is an analytical solution to the N-body problem – the Wang's solution (Wang 1991) – this solution remains useless for scientists. The reason for that is that the solution is expressed as a convergent infinite series, whose speed of convergence is too low. Numerical methods thus yield more reliable solutions than this analytic solution for the trajectories of N-body systems.

Second, we claim that analytical methods used to provide analytical solutions are sometimes too much sophisticated mathematical machinery for the problem at stake. In this case, analytical solutions, although available, are inadequate for the type of problem of which they are the exact solution. We discuss in particular the case of the simple pendulum in classical mechanics. This problem is often misconceived as being analytically intractable and therefore as necessarily requiring a numerical method to be solved (Gallant 2012, p. 70). Nevertheless, an analytical solution is available and numerical methods are, in principle, not required to solve the problem. However, to our knowledge, the analytical solution is never used in practice. This results from the fact that the analytical solution is a mathematical function expressed with Jacobi special functions, which require sophisticated mathematics, too much sophisticated given the simplicity of the simple pendulum system. Thus it turns out preferable for scientists to use a numerical method to evaluate the behaviour of the simple pendulum system.

Third, we argue that analytical methods do not offer a systematic approach for solving equations of different kinds as numerical methods do. Solving equations analytically seems to require specific mathematical techniques depending on the equation at stake. For example, the way to solve linear differential equations is not the same as to solve non-linear equations; similarly the way to solve homogenous equations is not the same as to solve non-homogenous equations. It results that solving analytical equations is sensitive to little modifications in equations: a little change in an analytically solved equation may lead to an analytically intractable equation. On contrary, using numerical methods is a more generic practice of solving equations. The same numerical method can be used to solve very different kinds of equations. This property is very important in the context of scientific modelling, which may require to modify the equations by adding new parameters and variables or by changing the power of an exponent.

In a nutshell, we argue that, while preferring ease to get numbers over exactness is a matter of pragmatic choice, such a choice is based on very mathematical properties of analytical solutions that we identity and discuss. We thus conclude that the alleged superiority of analytical solutions must be mitigated.

References

Fillion, N. & Bangu, S (2015). Solutions in the Mathematical Sciences and Epistemic Hierarchies, Philosophy of Science. (forthcoming)

French, A. P. & Taylor, E. F. (1998). An Introduction to Quantum Physics, Cheltenham : Stanley Thomas.

Gallant, J. (2012). Doing Physics with Scientific Notebook. A Solving Problem Approach, Chichester : Wiley.

Humphreys, P. (2004). Extending Ourselves: Computational Science, Empiricism, and Scientific Method, New-York: Oxford University Press.

Wang, Q.-D. (1991). The Global Solution of the N-Body Problem, Celestial Mechanics and Dynamics Astronomy, 50, p. 73-88.

Marina Baldissera. In what sense is uncertainty intrinsic to climate science?

Recent philosophical literature (Parker (2010), Frigg et al (2014), Stainforth et al. (2007)) has been focussing on different kinds of uncertainty that arise in climate science. The three main sources of uncertainty have been found to be initial condition uncertainty, parameter uncertainty and structural uncertainty. Structural model uncertainty is receiving particular interest due to the difficulties that arise in representing the relevant variables that drive atmospheric phenomena. Climate scientists have in fact produced a wide variety of models and agree that their models all contain highly idealizing assumptions. What makes things even worse, is that different models have different if not contradicting modeling assumptions (Oreskes et al, 1994).

The interest that these issues has sparked in philosophers prompts the question about whether structural model uncertainty is intrinsic to climate science or whether it is endemic to most sciences: is there something about modeling the atmosphere, or the climate in general, that makes the scientific enterprise of modeling uncertain in some intrinsic way that is in principle and in practice different from other branches of science?

In this paper, argue that structural model uncertainty is not intrinsic to climate science, or at least it is premature to make such a claim. There certainly are modeling challenges that make modeling the climate a particularly daunting task, yet climate science shares general difficulties found in other sciences. I will argue for this claim in two stages: in the first part of the paper I will show that climate science has developed in a particular way, and that there is a so-called 'gap' between theoretical understanding of the climate and the models that are used to obtain projections for the future climate (Held (2005), Knutti (2008)). In the second part I will draw some analogies between epistemic challenges and its consequences for ontology in biological and climate sciences. This is partly motivated by the fact that some philosophers have already made such a connection (see, e.g. Lloyd 2010) and that climate scientists hope for a model complexity hierarchy similar to the one (allegedly) naturally presented by biology (Held, 2005).

In the first section I argue that the gap is partly responsible for the fact that uncertainty is a major issue in using climate projections for policy-making. A lack of hierarchy of complexity in climate modeling impedes the development of a systematic model construction strategy (Held 2005). I claim that the lack of hierarchy is a consequence of not being able to identify variables, parameters and boundaries of various components in the atmosphere, and the interaction between these components at different scales. As general circulation models (GCMs) become more sophisticated, they add fine-grained details in an unsystematic way. This impedes the individuation of general physical principles that are responsible for the organization of various atmospheric structures at different scales. In light of this, it is therefore premature to conclude that structural uncertainty is intrinsic to climate science.

In the second part of the paper, I draw analogies between uncertainty found in climate science and uncertainty found in other sciences. I will show that there is a difference between modeling the atmosphere and modeling biological systems as suggested by Held (2005), but this difference is not decisive in determining whether there are different kinds of uncertainty arising in climate models and biological models. To make this analogy, I will draw upon the work of Levins (1966) to emphasize the epistemological similarities and Wimsatt (2007) to emphasize the ontological similarities between biological and climate modeling.

Levins identifies various conceptual and computational issues that arise in population biology. As in climate science, the issues arise because of the complexity of the target system. To be able to construct a model, scientists need to find a way of constructing a model that both allows for mathematical tractability and meaningful representation of the processes involved in the target system. I argue that the epistemic compromises faced by population biologists as described by Levins are of the same kind as the compromises faced by climate scientists: population biologists are concerned with finding the relevant level of organization such that a simplified model can be constructed. Similarly, climate scientists face epistemic difficulties when identifying the system’s boundaries its parameters and its dynamical variables (Emanuel 1986). Therefore, structural uncertainty that emerges from epistemic practices may be quantitatively but not qualitatively different in the two sciences.

Wimsatt identifies two different kinds of complexity that can be attributed to natural systems. These are descriptive and interactional complexity. The former captures how different theoretical perspectives of one given system capture spatiotemporal boundaries of the components of the target system. The latter is partly dependent on the former, and is a measure of the complexity the various causal component of the system, once the system is individuated. Wimsatt claims that in biology, individual subsystems `cross boundaries between theoretical perspectives and their decompositions' (2007, p. 185). This, he claims, is characteristic of complex biological systems. I argue that similarly, there are no clear boundaries of systems described by different theoretical perspectives in climate sciences. For example, the way a hurricane is modeled at one scale (e.g. for weather prediction), is not integrated with the way hurricanes are modeled in large scale GCMs (Davis and Emanuel, 1991). This lack of integration is due to a lack of agreement of how large scale thermodynamics relates to smaller scale fluid dynamics.

If ontological complexity is similar in biological and climate models, structural uncertainty in modeling either of these natural systems will share the same issues.

I conclude by claiming that indeed, the epistemological and ontological issues posed by the biological sciences are similar to the ones in climate science. This suggests that structural uncertainty is not intrinsic to climate science, but endemic to many modeling practices in the sciences.

Brice Bantegnie. A Shift in Focus: From Mental States to Mental Capacities

Metaphysical inquiries about the mind have almost exclusively taken mental states as their object. Among mental states are usually listed propositional attitudes (belief that p, desire that p,...), perceptual states (seeing, hearing,...), sensory states (pain, kinesthetic experience), emotions and many others.

The main question in the metaphysics of mind is then the question of the nature of these mental states: are they identical or not to physical states and if there aren't what is the relationship between physical and mental states. A simple look at the table of content of a handbook in philosophy of mind is enough evidence for this claim (Braddon-Mitchell and Jackson 1996, Clark 2001, Heil 2004, Bermudez 2005, Kim 2006). Mental states, if there are such things, belong to the explanans of cognitive science, that is, theses concepts play a central role in our psychological explanations. There are, however, other concepts which play a very important role in cognitive science, namely, concepts of mental capacities and behavioral capacities. Humans have multiple mental capacities (the capacity to see, to remember something, to recognize solve mathematical problems, to ascribe mental states to others, ...) and behavioral capacities (the capacity to walk, to grasp objects, to coordinate their movements with others...) These concepts of mental capacities do not belong to the explanans of psychological explanations, they are their explanandum. Such capacities have attracted very little attention. The quasi-exclusive focus on mental states seems to suggests that mental and behavioral capacities are unproblematic entities, but if this is the underlying assumption, then it is deeply mistaken. My aim in this paper is not only to devote some time to a neglected subject but to show that this neglect is very problematic in that it leads us to set aside valuable philosophical positions and to lose ourselves in useless debates. I take eliminative materialism as an example of a position which should be reevaluated and the debate on the individuation of the sense as an example of a sterile debate.

My paper is in three parts. In the first part I present my account of mental capacities. I argue that mental capacities are dispositional exemplified by individuals as a whole, this exemplification being grounded in the exemplification of dispositions by parts of the individual. The structure of the mechanisms in question is the categorical basis of these sub-capacities (Craver 2007). These mechanisms have states. To the extent that these states corresponds to the mental states posited by folk psychology, the mental part of folk psychology is vindicated. In the second part of the paper, I show that this might not be the case. I review the different capacities which are being investigated in cognitive psychology by taking a look at the table of content of textbooks in cognitive science. I show that a greater attention payed to the work of psychologists leads to a great diversity of capacities, that explanations of these capacities in terms of faculties, that is, modules have been prominent in recent cognitive science and that this should lead us to a reassessment of the debate between the eliminative materialists, who contend that the mental states posited by folk psychology don't exist (Stich 83, Churchland 89) and those who argue for the existence of mental states in favor of the eliminative materialism. In the third part of the paper, I argue that the good criterion of individuation can lead us to postulate a very high number of sensory modalities. There is already an extensive literature on these question of the individuation of mental capacities in the philosophy of perception. Indeed, sensory modalities are capacities and a number of criteria have been put forward to individuate the senses (Macpherson 2011) (Biggs, Matthen, Stokes 2014). In the literature on the distinction of the senses, the main criteria which have been put forward are the criteria of the properties that are represented (content criterion), the specific qualitative character of the experience (qualia criterion), the physical property of the stimulus (stimulus criterion) and the organ (organ criterion). To the extent that it limits itself to the senses, this literature is parochial and as a consequence misses the mark. Moreover, when we try to put these criteria to work in order to individuate mental and behavioral capacities in general, then the only criterion which remains plausible is the organ criterion. The consequence is that it is very likely that there is a very wide variety of senses.

Marie Barnett. Reasons and Conditional Preferences

In her highly influential work on social norms, Cristina Bicchieri has shown that people sometimes have conditional preferences. Specifically, Bicchieri has shown that a person's preference to follow some behavioral rule may be conditional on her expectations regarding the behavior and normative expectations of others. This concept of conditionality is central to Bicchieri's definition of social norms.

According to Bicchieri, “[a behavioral rule] is a social norm in a population...if there exists a sufficiently large subset [of the population] such that, for each individual i [within that subset]:

Contingency: i knows that a rule R exists and applies to situations of type S;

Conditional preference: i prefers to conform to R in situations of type S on the condition that:

(a) Empirical expectations: i believes that a sufficiently large subset of [the population] conforms to R in situations of type S;

and either

(b) Normative expectations: i believes that a sufficiently large subset of [the population] expects i to conform to R in situation of type S;

or

(b’) Normative expectations with sanctions: i believes that a sufficiently large subset of [the population] expects i to conform to R in situations of type S, prefers i to conform, and may sanction behavior.” (Bicchieri, 11)

In this paper, I argue that there are a number of reasons why an individual’s behavior may be conditional on the behavior and expectations of others. For example, a person’s behavior may be conditional on the behavior and expectations of others if:

1. The behavior and expectations of others are taken as evidence of:

a) A moral rule

b) A legal rule

c) A social rule (e.g. a rule of etiquette)

d) Another kind of rule (e.g. a rule in a game)

Any of which might be associated with potential sanctions, which in turn can be:

i) social* (e.g. expressions of disapproval)

ii) non-social (e.g. the refusal of some tangible benefit, like money)

*I use the word “social” in a special sense, to indicate the emotional aspects of social interactions, abstracted away from the purely practical elements. However, in reality these things are almost always closely intertwined (e.g. I do not want my neighbor to disapprove of me both because the loss of neighborly friendliness/sense of community is emotionally upsetting, and because I know this disapproval will mean I am less likely to receive help with practical problems in the future). In theory, it is possible to consider both motivational elements separately; in reality, it may not be possible to pull them apart.

2. The behavior and expectations of others are taken to determine:

a) The morally correct action (i.e. if some moral rule makes reference to the behavior and expectations of others; e.g. an obligation to keep your promise only when the promisee – perhaps among others – expects you to do so, and you are within a group where promises are meaningful because they are generally kept), where compliance with the moral rule may be motivated by:

i) internal factors (moral determination), and/or

ii) external factors (sanctions, in this or some other life)

b) The social rule (e.g. when the behavior and expectations of others determine “the way we do things around here”), where compliance with the social rule may be motivated by:

i) internal factors (a desire to do things in the socially accepted way), and/or

ii) external factors (sanctions, which may be social or non-social)

3. The behavior and expectations of others motivate compliance with a pre-existing rule (e.g. if a person only wishes to be moral if others do so and expect her to do so, if she only wishes to follow the rules of a game when others do so and expect her to do so, etc.).

4. The behavior and expectations of others affect the person's motivation to comply with some conflicting rule in one of the above ways. For example: Assume only two behaviors, A and B, are possible within some situation. If behavior A is a social norm, compliance with which is motivated by the threat of social sanctions (where these sanctions are taken to be the worst possible outcome, e.g. because the approval of the community is very highly valued), while behavior B is motivated by something else (moral beliefs, desire for money, etc.), the presence of both behaviors will track people's beliefs about the behavior and expectations of others, even though these factors are only directly influencing behavior A.

I believe that it is extremely important to understand not only that people have socially conditional preferences (i.e. preferences which are conditional on the behavior and expectations of others), but also why they have such preferences. That is, I argue that it is important to understand the reasons that support this kind of conditionality, and to develop experiments which will allow us to distinguish between the different kinds of reasons outlined above. This is essential because, depending on which factors motivate people's conditional preferences, different interventions may be effective. For example, if people are taking the behavior and expectations of others as evidence of some rule, the testimony of an authority figure may be sufficient for behavioral change. Alternatively, if people's conditional preferences are motivated by fear of tangible, non-social sanctions, we may wish to consider compensating those who agree to stop following an undesirable norm. In this paper, I argue that the development of methods of distinguishing between the different kinds of reasons which might produce conditional preferences is a natural next step for social scientists involved in the study of social norms.

References

Bicchieri, Cristina. The Grammar of Society: The Nature and Dynamics of Social Norms. Cambridge: Cambridge University Press, 2006.

Michael Baumgartner and Lorenzo Casini. Establishing Constitutional Relations, in Theory and in Practice

It is a popular maxim in recent debates about mechanistic explanation that a powerful strategy to explain the upper level behavior P of some system S consists in pinpointing the lower level mechanism that constitutes P(S) (Glennan 2002; Bechtel and Abrahamsen 2005; Craver 2007). This raises the methodological follow-up question as to how mechanisms are best identified, i.e. how those of S's spatiotemporal parts X = {X1,...,Xn} are singled out, whose causal activities A(X) = {A1(X1),...,An(Xn)} are constitutively relevant to P(S). According to a prominent answer due to Craver (2007), constitutional relations are experimentally uncovered along roughly the same lines as causal relations—notwithstanding the fact that constitution and causation are very different relations (Craver and Bechtel 2007).

Since the time of Mill (1843), one of the dominant experimental approaches to uncovering causal relations, influentially systematized by Woodward (2003), consists in intervening on causes (in controlled environments) to change their effects. Craver (2007) argues that the same basic idea—with a mutuality tweak—applies to discovering constitutional relations. Subject to his mutual manipulability account of constitution (MM), the behavior Ai(Xi) of a spatiotemporal part Xi of S is a constituent of P(S) iff it is possible to 'ideally' intervene—in the sense of Woodward (2003, 98)—on Ai(Xi) such that P(S) changes, and on P(S) such that Ai(Xi) changes (Craver 2007, 153). Identifying constitutional relations along the lines of MM, for Craver, is not only a theoretical proposal but a faithful reconstruction of scientific practice.

By drawing on a recent result of Baumgartner and Gebharter (2015), the first part of this paper shows that MM does not ground an adequate methodology for constitutional discovery. In short, the reason is that the idealized experiments required by MM are unrealizable in principle, for upper level phenomena and their constituent mechanisms are so tightly coupled that they can only be manipulated with a fat-hand, i.e. via common causes. Furthermore, less rigorous but realizable experimental set-ups systematically underdetermine the inference to constitutional relations, due to the (non-ideal) fat-handed nature of relevant manipulations. In sum, while there exist experimental designs that, given compliance with required assumptions about unmeasured background influences, conclusively establish the existence of causal relations, no such experimental designs can possibly exist for the inference to constitutional relations. Therefore, the inference to constitutional relations cannot proceed along the lines of the inference to causal relations. If scientists were to follow MM's prescriptions, their reasoning would be fallacious. Hence, if we grant that their reasoning is not fallacious, it must be reconstructed differently.

Inspired by suggestions from Simon (1962) and Wimsatt (1997), the second part of the paper draws on recent research in neuroscience to develop an 'abductivist' alternative to MM. Neuroscientists increasingly rely on network theory (Newman 2006) to split networks of unit activations in the brain, and patterns of co-activation between them, into distinct component modules, in order to map different cognitive/behavioral phenomena onto the different modules (Nelson et al. 2010; Meunier et al. 2009). Resulting decompositions are considered adequate if they account for the phenomena under investigation to a high degree of accuracy, and without redundancies, i.e. no element of the decomposition may be removed without a loss in accuracy. Adequacy of a decomposition guarantees constitutional relevance. Once an adequate decomposition has been recovered, its robustness is tested by varying the number or size of the parcellated units whose activation is being analyzed, by adding or deleting some of their putative causal connections, by modifying the descriptive grain of units and connections, and by redescribing the phenomenon, or phenomena, to be explained. When the decomposition is shown to be robust, the set of constituents is taken to provide the optimal, or most relevant, mechanistic decomposition. This procedure, we argue, is best reconstructed as an attempt to give an empirically accurate, redundancy-free, simplest and most unifying account of the phenomena—that is, to provide a maximally powerful explanation.

We propose an approach to constitutional discovery that generalizes this pattern of reasoning. Constitutional relations are established by way of abductive inferences. More concretely, the constituents of a mechanism for an upper level behavior P(S) are recovered by decomposing the corresponding system S into a set of proper spatiotemporal parts X whose causal activities A(X) constitute P(S). This goal is accomplished, we contend, if the decomposition satisfies the following constraints:

1. Accuracy. The set of activities A(X) are sufficient to deduce P(S).

2. Coupling. The behavior P(S) and the elements of A(X) are so tightly coupled that: (i) all causes of P(S) are common causes of P(S) and some Ai(Xi); (ii) every Ai(Xi) has at least one 'non-structure-altering' cause that is a common cause of Ai(Xi) and P(S).

3. No de-coupling. P(S) and P(X) resist de-coupling across all expansions of the variable sets P(S) and A(X).

Against the background of our proposal, the role of top-down and bottom-up manipulations by fat-handed interventions on P(S) and A(X) is not the one depicted in MM. Contrary to MM, successful combinations of top-down and bottom-up experiments are never sufficient to warrant the inference to constitutional relations. Rather, they are a means to establish the coupling of P(S) and A(X)—in line with (2)—and to test whether this coupling can be broken—in line with (3). The sufficiency of A(X) for P(S)—in line with (1)—together with evidence of coupling and persistent failure of de-coupling attempts are best explained by introducing a constitutional dependency between P(S) and A(X).

(1) to (3) suffice to identify an adequate set X of constituents. However, there may be multiple such sets: X, Y, Z, etc. Then, the pragmatic issue arises as to how to choose between them. In neuroscience, scientists add to constraints (1) to (3) additional desiderata, for instance:

4. Hubness. The decomposition of P(S) is (i) robust across 'structure-altering' bottom-level manipulations on 'peripheral' nodes and (ii) sensitive to 'structure-altering' bottom-level manipulations on 'hub' nodes.

5. Rescaling. The decomposition of P(S) is robust across 'structure-altering' bottom-level manipulations that modify the size of the units, thereby changing number and kind of connections.

6. Unification. The decomposition of P(S) is coherent with decompositions of further phenomena Q(T), R(U), etc. into constituents that comply with (1) to (3) and overlap with X.

Whereas (1) to (3) are normative guidelines that justify the inference to constitution, (4) to (6) (and possibly other desiderata) act as pragmatic guidelines that help select one set of constituents as most explanatory.

Benjamin Bewersdorf. Conceptual Learning and Bayesian Epistemology

When a child is born, it does not know whether the moon is made of cheese or stone. It does not know this for two reasons. The first reason is that it lacks factual knowledge about the moon. The second reason is that it does not even possess the conceptual resources necessary to form beliefs about the moon. When the child grows up, it will both learn the relevant facts and the necessary concepts. Bayesian epistemology offers a successful and well explored account of factual learning. Conceptual learning, however, is not as thoroughly investigated from a Bayesian perspective (see however Williamson (2003), Romeijn (2006), Romeijn and Wenmackers (forthcoming)).

According to the Bayesian, the belief state of a rational agent can be represented by a probability distribution on an algebra of propositions. The propositions in this algebra constitute all the propositions towards which the agent has an opinion. The probabilities assigned to these propositions represent the degrees to which the agent believes that the propositions obtain. Factual learning can be understood as a change in the agent's subjective probability distribution.

When an agent learns a new concept she becomes able to form beliefs that involve this concept. This means that the number of propositions towards which the agent has an opinion is increased. In the Bayesian framework this amounts to an extension of the algebra on which the agent's subjective probability distribution is defined. However, a mere extension of the agent's algebra cannot be all there is to learning a new concept. Otherwise there would be no difference between learning the concept of a bird and learning the concept of a fish. Furthermore, we would not say of someone that she has learned the concept of a bird unless she also developed some understanding of what it means to be a bird. Thus, learning a concept requires the agent to acquire information (Williamson (2003) agrees that conceptual learning involves receiving information, while Huber (2009, pp. 23-4) disagrees).

There are different kinds of information that an agent might have to acquire in order to count as having learned a particular concept. In this paper I focus on information about the inferential relations between concepts. As has been argued by Pollock (1989, ch. 4 and 5) and Pollock and Cruz (1999, pp. 147-

150), concepts are individuated by their inferential relations. This means that learning a particular concept requires the acceptance of particular inferential relations. Failing to do so amounts to failing to learn this particular concept. In a Bayesian framework such inferential relations can be represented by a set $\mathbb{C}$ of constraints on the conditional degrees of belief of the agent. Such constraints can be more or less specific. In the extreme case they completely determine what degrees of belief the agent should adopt after having learned a concept. In such a case the following seems to be a natural suggestion for a Bayesian account of conceptual learning: the agent should extend her algebra to include the propositions introduced by the new concept and set her degrees of belief to the values implied by $\mathbb{C}$.

For cases in which $\mathbb{C}$ does not completely specify the new degrees of belief Jon Williamson (2003) has proposed the following account of conceptual learning. The agent should first determine the probability distributions on the extended algebra which comply to $\mathbb{C}$ and are most similar to the previous probability distribution (where similarity is defined by cross entropy on the old algebra). She should then choose among these the one which distributes the probability mass most evenly (where even distribution is defined in terms of entropy).

Both of these accounts assume that $\mathbb{C}$ imposes constraints on the new degrees of belief of the agent. However, if we think of $\mathbb{C}$ as representing the inferential relations that individuate the concepts in question, this assumption causes several severe problems. In this case it can be shown that both accounts violate a version of the principle of language invariance and the principle of commutativity, allow for cases in which learning facts has no influence on the agent's degrees of belief and in which learning facts can lead to the agent losing conceptual information.

I argue that these problems can be addressed by a simple modification of these accounts. Instead of understanding the constraints imposed by learning a concept as constraints on the new belief state of the agent, we should understand them as constraints on the initial or a priori belief state of the agent. That is the belief state the agent would be in if she had only received conceptual information and no factual information. Factual learning should then be represented by conditionalizing this modified a priori belief state on the total evidence of the agent. I show that doing so resolves the problems introduced above.

References

Huber, F. (2009): Beliefs and Degrees of Belief, in: Huber, F.; Schmidt-Petri, C. (eds.): Degrees of Belief, Springer

Pollock, J. L.; Cruz, J. (1999): Contemporary Theories of Knowledge, second edition, Rowman & Littlefield Publishers

Pollock, J. L. (1989): How to Build a Person. A Prolegomenon, MIT Press

Romeijn, J.-W. (2006): Theory Change and Bayesian Statistical Inference, Philosophy of Science S72:5, pp. 1174-1186

Wenmackers, S.; Romeijn, J.-W. (forthcoming): A New Theory about Old Evidence, Synthese

Williamson, J. (2003): Bayesianism and Language Change, Journal of Logic, Language and Information 12, pp. 53-97

Harjit Bhogal. Three Dimensions of Explantory Goodness

We often judge one explanation of a proposition to be better than another explanation of the same proposition. Such judgments occur in situations where we have two complementary explanations for an event — e.g. when we are comparing fundamental physical and genetic explanations of an event —

and in situations where we have two competing explanation for an event — e.g. when we are comparing explanations in a process of inference to the best explanation. In this paper I give an account of this notion of explanatory goodness.

In particular, I given an account of explanatory goodness which grounds the goodness of a particular explanation in terms of the pattern of explanatory facts in other possible worlds. I claim that the there are three dimensions of explanatory goodness which I call PRECISION, ROBUSTNESS and CHANCE.

Here is a case that motivates PRECISION as a dimension of goodness. Consider this statistical mechanical explanation:

Explanandum: A particular ice cube melts. Explanans I took that ice cube out of the freezer and dropped it in warm water.

This is generally taken to be a very good explanation. The reason is that nearly all of the ways the explanans could be true explain (deterministically if we are assuming classical statistical mechanics) the explanandum. Discussions of such SM explanations make a great deal of this fact — the overwhelming majority of ways the microstate could realize the explanans lead to the explanandum; there are a negligible number of `bad cases'.

Conversely, consider explanation like this:

Explanandum: A particular ice cube grows. Explanans I took that ice cube out of the freezer and dropped it in warm water.

Such an explanation seems bad because only very few of the ways the explanans could be true would explain the explanandum.

An explanation scores well on PRECISION if most of the ways that the explanans could be true explain the explandum. I'll give a more formal account of PRECISION soon.

Here is a case that motivates PRECISION as a dimension of goodness. Consider these two competing explanations.

Explanandum: A particular species recently became extinct. Explanans 1: A predator was introduced into the environment and quickly wiped out the species. Explanans 2: The initial conditions of the universe and the deterministic laws imply that the species would become extinct at the time it did.

Explanation 1 seems clearly better. I consider some accounts of this betterness (in particular Jackson and Pettit's account of `modally comparative information') and argue that the reason that the first explanation is better is that the explanation holds in more of the close worlds where the explanadum is true. The explanation is not modally fragile and thus scores well on ROBUSTNESS.

Now to give a more formal account of PRECISION and ROBUSTNESS.

PRECISION and ROBUSTNESS can both be characterized by introducing the notion of counterfactual probability.

S $\rightarrow$ (x) T is true iff x of the close S-worlds are worlds where T

Here we can say x is the counterfactual probability of T given S. x gives the proportion of close-S worlds where T is true. That is, to x is given by counting the number of T worlds and dividing by the number of S worlds. I should say something about the characterization of closeness here. Let us say that, in the context of considering an explanation of T from S, the close S-worlds are all the nomically possible S-worlds that are consistent with the background conditions, M, that we are assuming when we are giving the explanation. (Clearly we will needs a measure over worlds here, I suggest that the standard measure used in statistical mechanics is the one we want.)

Here's how to characterize PRECISION in terms of counterfactual probability. Consider:

A → (x) A explains B

Here, x is the proportion of nomically possible worlds that are consistent with A and M where A explains B. The PRECISION of an explanation of B from A is higher if x is higher.

Let's move on to ROBUSTNESS. Consider:

B → (x) A explains B

Again, closeness here is characterized as worlds that are nomically possible and consistent with the background conditions.

An explanation scores highly on ROBUSTNESS if in most of the possible worlds where the B holds A explains B.

Together, PRECISION and ROBUSTNESS imply that an explanans should be correctly `targeted' at an explanandum. If the explanans holds in lots of worlds where the explanandum does not then the explanation will score low on PRECISION. If there are many worlds where the explanandum holds but the explanans does not then the explanation will score low on ROBUSTNESS.

So far we have been ignoring probabilistic explanations. The third dimension CHANCE is relevant when we are taking them into consideration. An explanation scores better on CHANCE if the explanans implies that the explanandum has a higher objective probability. (Though there are some complexities about the interaction between CHANCE and PRECISION.)

I discuss a case which shows that chance and counterfactual probabilities are very different types of probabilities. There is no tight connection between them. In fact, I argue that chance and counterfactual probabilities have very different features. In particular, they play very different roles in explanation; counterfactual probabilities are involved in what Skow calls almost necessity explanations. What's more, chances are law governed magnitudes in a way that counterfactual probabilities are not.

Seeing these differences not only sheds light on explanatory goodness but also on the nature of probabilities, particularly in deterministic worlds. In particular, once we have recognized that there are two different senses of non-epistemic probability that play a role explanations and have different connections to laws then we have more options for how we can construe the nature of probabilities in deterministic worlds. In particular, we can reject deterministic chance whilst still accepting that there is an objective, explanatory probability — deterministic counterfactual probability.

Christopher Blunt. How to Create False Positives and Influence People: cohort multiple RCTs and the Grades of Recommendation

In this paper, I demonstrate that there are trial methodologies which are (at least superficially) randomised controlled trials which can be relied upon to generate false positives in the absence of an effective treatment. I show that our current protocols for evaluating evidence in health policy and clinical practice are susceptible to serious exploitation by trials using these methods. I then present recommendations for defending against this exploitation on the basis of a more rigorous of evidence, which acknowledges the potential for an asymmetry between evidence for and against the effectiveness of a treatment.

First, I argue that the trial methodology known as "cohort multiple randomised controlled trials" or cmRCTs is in effect a false positive generator. cmRCTs were suggested in the BMJ by Relton et al. [1] in 2010 as a methodology which combines the best of both randomised controlled trials and observational studies. I argue that, in fact, such trials are not properly controlled, and systematically favour the experimental treatment. I claim that for any non-harmful intervention, we should expect cmRCT results to support the experimental treatment, generating a false positive if the treatment is not effective. In particular, cmRCTs have been used to generate positive results interpreted as demonstrating the efficacy of complementary and alternative medicines, such as homeopathy (see [1-5]).

I then present an influential current approach to the appraisal of evidence in medicine and health policy—'grades of recommendation'. Grades of recommendation often accompany hierarchies of evidence in policy documents and in guidance to clinicians and policymakers. They allow users to appraise the strength with which they should recommend a treatment on the basis of the evidence available in favour of it. Most such documents claim that the strongest recommendations are justified when there are RCTs or meta-analyses of RCTs which support the effectiveness of the treatment. This approach has been widely implemented by governmental bodies such as NICE in the UK [6,7] and the Australian NHMRC [8], guideline producers such as the Scottish Intercollegiate Guideline Network (SIGN) [9], the Canadian Task Force on the Periodic Health Examination [10] and the US Preventive Services Taskforce [11], and by leading proponents of Evidence-Based Medicine (EBM) (e.g. [12,13]).

I show that grades of recommendation are open to exploitation by false positive generating methodologies such as cmRCTs. Most grades of recommendation approaches only consider positive evidence and fail to appreciate the significance of negative findings. I argue that these approaches allow cherry-picking of positive findings, as well as deliberate research design choices which favour positive outcomes, such as the use of cmRCT methodology—in particular in cmRCTs, the use of 'Zelen consent' or 'post-randomised consent' [14,15]. The case of cmRCTs should give serious cause for concern about basing treatment guidelines or clinical practice upon grades of recommendation.

Finally, I present a set of proposals to improve the current approach to the appraisal of research evidence in clinical practice and health policy in response to these issues. I argue for consideration of negative as well as positive findings, and a greater focus upon particular methodological choices (e.g. blinding, choice of null hypothesis, control intervention, consent protocols) in addition to overall methodology. I present an approach which uses different criteria to evaluate the evidence for vs. against the effectiveness of an intervention, arguing that there is an asymmetry in the quality and strength of different methodologies in showing that interventions are or are not effective.

References

1. Relton, Clare, et al. (2010) "Rethinking pragmatic randomised controlled trials: introducing the "cohort multiple randomised controlled trial" design". BMJ, 340.

2. Relton, C. (2009) A new design for pragmatic RCTs: a "patient cohort" RCT of treatment by a homeopath for menopausal hot flushes. University of Sheffield: [PhD Thesis] ISRCTN 0287542.

3. Relton, Clare, et al. (2011) "South Yorkshire Cohort: a'cohort trials facility'study of health and weight-Protocol for the recruitment phase". BMC Public Health, 11(1), 640.

4. Relton, C, O'Cathain, A & Nicholl, J. (2012) "A pilot 'cohort multiple randomised controlled trial'of treatment by a homeopath for women with menopausal hot flushes". Contemporary Clinical Trials, 33(5), 853-859.

5. Viksveen, Petter & Relton, Clare. (2014) "Depression treated by homeopaths: a study protocol for a pragmatic cohort multiple randomised controlled trial". Homeopathy, 103(2), 147-152.

6. National Institute for Health and Clinical Excellence (NICE). (2004) NICE: Guideline Development Methods (Vol. Vol.7: Reviewing and Grading the Evidence).

7. National Institute for Health and Clinical Excellence (NICE). (2011) CG117: Tuberculosis: clinical diagnosis and management of tuberculosis, and measures for its prevention and control. London: RCP.

8. Australian National Health and Medical Research Council (ANHMRC). (1999) A Guide to the Development, Implementation and Evaluation of Clinical Practice Guidelines. Commonwealth of

Australia: available at: http://www.health.gov.au/nhmrc/publicat/synopses/cp30syn.html.

9. Scottish Intercollegiate Guidelines Network (SIGN). (2008) SIGN 50: A Guideline Developer's Handbook. available at: www.sign.ac.uk.

10. Canadian Task Force on the Periodic Health Examination. (2014) Procedure Manual. http://canadiantaskforce.ca/methods/.

11. U.S. Preventive Services Task Force. (2008) U.S. Preventive Services Procedure Manual. AHRQ Publication No. 08-05118-EF.

12. Howick, J., et al. (2011) “The Oxford 2011 Levels of Evidence”. Oxford Centre for Evidence-Based Medicine, available at www.cebm.net.

13. GRADE Working Group. (2004) “Grading quality of evidence and strength of recommendations”. British Medical Journal, 328, 1490.

14. Zelen, Marvin. (1979) “A new design for randomized clinical trials”. N Engl J Med, 300(22), 1242-1245.

15. Adamson, Joy, et al. (2006) “Review of randomised trials using the post-randomised consent (Zelen's) design”. Contemporary Clinical Trials, 27(4), 305-319.

Seamus Bradley, Karim Thébault and Alexander Reutlinger. Modelling inequality

Econophysics is a cross-disciplinary research field that applies models and modelling techniques from statistical physics to economic systems. Methodologically, econophysics is supposed to differ from conventional economic practice in that it uses the `paradigms and tools’ of statistical physics. Econophysicists hold that mainstream economics suffers from a number of defects that they would like to correct. Most importantly both the core principles and models of mainstream economic theory are argued to lack the evidential support of real economic data. Econophysicists suppose that models and modelling techniques drawn from an experimentally focused and mathematically sophisticated science such as statistical physics will give new, and more reliable insights. Not surprisingly, not all mainstream economists agree with such a dismal view of their science. Authors offer a sharp critique of econophysics on the grounds of some practitioners: i) redoing work which has been done within economics; ii) ignoring rigorous and robust statistical methodology; iii) assuming universal empirical regularities where there are none; and iv) using modelling techniques that are in certain senses inherently problematic or illegitimate, above all that econophysical models suffer from completely unjustifiable strong idealisations. Considering this final line of criticism in the context of econophysics models of income inequality is the main focus of this paper.

The hallmark of econophysics models is their success in capturing certain `stylised facts’ found within economic systems. Simple physics-inspired models can reproduce important distributional features of economic systems, such as the scale freedom of price fluctuations in financial markets, or the `power-law tail’ of the distribution of monetary income or wealth in populations. Perplexingly, recovery of the income (or wealth) `stylised fact’ can be achieved within extremely simple, and heavy idealised, econophysics models of monetary exchange. Drawing on analogies with statistical mechanics, these `kinetic exchange’ models of income or wealth distributions model economic agents as zero-intelligence particles who bump into each other and exchange money between them at random – in many respects just like the molecules in a gas. Despite painting an idealised picture of economic interactions that is quite far removed from reality, these simple statistical physics-inspired models are remarkably successful at capturing the broad features of the distribution of income within populations. The critique of kinetic exchange models found in the literature is primarily a methodological attack regarding the idealisations involved in the models: it is not the accuracy of these models in recovering real data that is in question. Rather, kinetic exchange type models for inequality are argued to be inherently problematic or illegitimate on the grounds that their treatment of (for example) production, income and transactions is in conflict with `economic reality'. In our paper we will assess the warrant of the criticisms drawing both upon the philosophical literature on modelling and idealisation, and upon the notion of a 'maximum entropy explanation'.

One important aspect of our paper, is a comparison between the various idealisations made in the econophysics model, and in the statistical physics model that inspired it. We will highlight how the idealisations can be justified in the two contexts of statistical mechanics and econophysics, and in so doing, rebut some of the criticisms aimed at this model in the literature. This serves as a platform to discuss some general lessons about the importance of background knowledge in justifying idealisations. More specifically, our analysis will involve a comparison between the respective idealisations relating to i) binary interactions; ii) conservation principles; and iii) the exchange dynamics. In each case the similarity in idealisation between the income and gas model will be contrasted with the difference in mode of justification. For instance, whereas in the gas case binary interactions are a legitimate approximation to certain density regime, for inequality, we argue that the binary interaction idealisation is only justified in terms of the model's retention of an explanatorily relevant factor (entropy maximisation) – for this reason we argue that the idealisation in question is a minimalist idealisation.

Although the idealisations involved in kinetic exchange models of inequality must be justified in different terms to those involved in kinetic exchange gas models, there can therefore still exist legitimate methods for their justification. To this end, we will consider in detail the foundations of explanations via maximisation of entropy, drawing upon recent work in the foundations of statistics and in philosophy. We will point to a number of conceptual problems relating to the employment of the notion of `entropy' within economics, and critically examine the putative justification of the idealisations found in kinetic exchange income models via entropy maximisation. Our paper will thus have both specific implications for the debate regarding kinetic exchange models of inequality, and wider implications for analysis of models and idealisation in econophysics and beyond (including but not limited to a critical discussion of understanding econophysical explanations as structural mechanistic explanation). Our hope is to offer guidance with regard to both the practice of modelling inequality, and the inequality of modelling practice.

Daniel Calder. Ramsey Reconsidered: Applying the job-descripton challenge to contemporary cognitive science

When considering whether a cognitive system deals in representations, it is important to examine the function of states or processes that are purportedly representing, and compare this with a properly representational function. If the entity under scrutiny performs a function that is not representational, then we ought not call it as such (even if a representational gloss is possible). This is William Ramsey's job-description challenge. With it, Ramsey hoped to show that while classical approaches to cognitive science do invoke properly representational states, contemporary methods do not, and researchers would make more fruitful contributions by eliminating representation-talk from their theorising (Ramsey, 2007).

However, several reviewers and commentators (Sprevak, 2011; Grush, 2008; Shagrir, 2012) have criticised Ramsey's limited analysis of contemporary science. They suggest that several paradigms we neglected – most notably forward model control theory and attractor dynamics. Furthermore, they argue that scientists using these techniques to model cognitive phenomena do employ a notion of representation that Ramsey permits – so called S-representations, a catch-all term for 'similarity representations' and 'structural representations' which both rely on a isomorphism relation between the representing state and its target. Thus S-representations are not to be forgotten in the history books of the field, but instead are alive and fertile posits being used today in our current best theories about cognition.

In this talk, I will defend Ramsey's eliminativism against such criticisms. First I will take on Oron Shagrir's detailed response to Ramsey, which appeals to an analysis of attractor networks in the oculomotor system taken to encode memories of eye-position as structural representations (Shagrir, 2012). Then I will discuss forward modelling, and its promising successor, predictive processing, which pose an important threat to the eliminativist case (Grush, 2008). The presence of a model of any sort suggests a structural similarity in the system that allows it to successfully regulate its task domain. Given the ubiquity of both model-based research and biologically plausible attractor dynamics, eliminativism must meet these challenges to remain tenable. Ramsey's job-description challenge affords a clear proving ground for both parties, which allows questions of content to be put to one side, instead encouraging philosophers to focus on the explanatory role of theoretical posits.

In his 2012 paper, Structural Representation and the Brain, Shagrir constructs an in-depth study of a neural integrator situated in the oculomotor system and characterises it as a recurrent neural network functioning as a memory. I will argue that Shagrir's evaluation is flawed, and that his conclusions do not follow from current theoretical understanding of the oculomotor mechanism. The alternative I will demonstrate shows that the integrator has direct control over eye position, excluding exceptional circumstances, such as having electrodes implanted to artificially stimulate motor neurons. The result is that the states Shagrir takes to be representational only trivially covary with their target, and are not used as representations, so they fail Ramsey's job-description challenge.

Forward models are elements of a control system that receive a copy of the current output signal, so that they can quickly construct an expectation of the state of the plant (the robot or system being controlled) and feed that prediction back to the controller. The forward model thus must be isomorphic to the plant itself, in order for its predictions to be useful, and functions as a representation because its explicit role in the system is to stand-in for the plant and be used as such by the control element. However, when the forward model approach is extended, as it is in the predictive processing paradigm, the representation relation no longer exists.

For advocates of predictive processing, the whole system embodies a forward model of its target, generating predictions and processing error. Though very new, this approach has scored important successes (e.g. Hinton, 2006) and is shaping up to be a powerful development upon the forward model approach. I will highlight an important change that has repercussions for the representational status of these expanded forward models, which revives the eliminativist project. When the model is no longer useful to some control system, and instead takes control itself, it no longer performs the important standing-in role. Rather, the system is called a model in virtue of a trivial, information-sensitivity, and not because of its function within a system. This kind of modeling fails the job-description challenge because states playing a representational function must be used by a system as surrogates for their target, but when the system itself is taken to be a model, there is no other process to make use of its potential as a model. So the system can simultaneously be a model, and not be functioning as a representation.

Whist I make a large concession to forward model theory, I hope my arguments will convince the conference that two contemporary and influential model-involving paradigms do not require representational posits, whether they refer to models based on attractor dynamics or models based on predictive processing. Eliminativism supported by Ramsey's job-description challenge can thus make some important claims over the posits of current cognitive science, and not be deterred by objections to Ramsey's original analysis.

References

Grush, R. (2008). Book review: Representation reconsidered. Notre Dame Philosophical Reviews.

Hinton, G. (2006). A fast learning algorithm for deep belief nets. Neural Computation, 18:1527-1554.

Ramsey, W. (2007). Representation Reconsidered. Cambridge University Press, New York.

Shagrir, O. (2012). Structural representations and the brain. British Journal of Philosophy of Science, 63:519–545.

Sprevak, M. (2011). Book review: Representation reconsidered. British Journal of Philosophy of Science, 62:669–675.

Kevin Coffey. Reconsidering Unconceived Alternatives: Prospects for Scientific Realism

This paper offers a novel realist defense against the so-called "problem of unconceived alternatives", and uses that defense to reassess the relationship between theory and evidence in foundational science. In his recent and influential book, Kyle Stanford attempts to undermine the epistemic status of scientific realism by appealing to the possibility of unconceived but scientifically serious and well-confirmed rivals to our best foundational theories. The existence of such unconceived alternatives would seem to show that the central claims of our best scientific theories are underdetermined by the available evidence, and thus that we should give up faith in the approximate truth of those theories. Although Stanford's "problem of unconceived alternatives" is both more sophisticated and more compelling than existing historical and underdetermination-based arguments against scientific realism, I argue in this paper that it fails to undermine the realist's epistemic warrant. Even if one concedes that there exist scientifically serious unconceived alternatives to contemporary foundational theories, Stanford historical induction fails to demonstrate that such unconceived alternatives are as well confirmed as existing theories. Indeed, there are good reasons to think otherwise—reasons that suggest ways the epistemic situation of contemporary science is importantly different from the epistemic situation encountered by successful science in the past. I then go on to argue that Stanford cannot avoid my critique without giving up precisely those features of his argument that were intended to make it more compelling than the traditional historical and underdetermination-based arguments against scientific realism. My paper thus aims to show where Stanford's argument goes wrong, and to suggest an illuminating but overlooked aspect of the theory—evidence relationship in foundational science.

The paper's argument is developed in four parts. Part one situates Stanford's problem of unconceived alternatives within the scientific realism debate, and identifies several important ways in which Stanford's argument is a good deal more sophisticated and plausible than traditional historical and underdetermination-based arguments against scientific realism. In part two I critically examine Stanford's historical induction for the conclusion that there are presently unconceived rivals that are as well confirmed by the available evidence as existing foundational scientific theories. Using specific examples from physics, I argue that the historical record doesn't support this claim. There are good reasons to think, for example, that the 17th century evidence available in support of classical dynamics didn't (and doesn't) support special relativity equally well. Rather, the historical record indicates that special relativity came to be accepted only on account of the emergence of new types of empirical evidence, and only for this reason came to supersede classical dynamics. Similar considerations hold, I argue, for a variety of other notable cases. As a consequence, the most that Stanford's historical argument establishes is that there are very likely unconceived rivals to contemporary science that are consistent with the available evidence. To avoid collapsing into a more traditional (and problematic) underdetermination argument, then, Stanford must show that there are currently unconceived rivals to foundational science that new types of evidence will (eventually) establish as better confirmed than existing theories. But, as I argue in part three, it's on precisely this issue that our current evidential situation appears different from that of past science, for there are good reasons to think that new types of evidence will not be readily forthcoming in future foundational science. This suggests an epistemically-relevant way in which the evidential situation of contemporary foundational science is quite different from that of successful foundational science in the past, and undermines the lessons Stanford wants to draw on the basis of unconceived alternatives. Nevertheless, I think there are reasons to be suspicious of scientific realism. I conclude in part four by discussing several sources of anti-realist doubt from within foundational science that haven't received the attention in the philosophical literature that they deserve.

Matteo Colombo and Jan Sprenger. Explanatory Value and Probabilistic Reasoning: An Empirical Study

The interplay of explanatory, causal, and probabilistic reasoning is tight and multidirectional. While the question of how judgments of explanatory value (should) inform probabilistic inference has been well studied within both psychology and psychology, e.g., in the literature on abductive inference, the related question of how probabilistic and causal information (should) affect judgments of explanatory value has received less attention.

One way to address this question is to begin with the hypothesis that explanation is "a two-tiered structure consisting of statistical relevance relations on one level and causal processes and interactions on the other" (Salmon 1997: 475-6). According to this hypothesis, explanatory value depends on the joint contribution of statistical relevance relations and causality: both factors are indispensable to explanatory value, which has also been stressed recently by the literature on probabilistic causation (e.g., Halpern and Pearl 2005; Hitchcock 2008).

In the present paper, we elucidate this hypothesis by addressing whether and under which circumstances judgments of explanatory value are associated with causal and probabilistic characteristics of a potential explanation.

To address these issues, we conducted two experimental studies. In both studies, experimental participants read well-constrained problem situations where information about statistical and causal relevance relations between an explanandum and a potential explanatory hypothesis was provided. Participants were asked to make a series of explanatory judgments along several dimensions, including judgments about the explanatory value of the hypothesis and its cognitive and causal relevance, but also about its plausibility, degree of confirmation and its logical relation to the evidence.

In the first study, we examined explanations for a certain type of event, where no alternative explanation was explicitly given, but many potential alternative explanations could be easily produced. We tested three specific hypotheses: (i) that judgements of explanatory value were reliably predicted by the prior subjective credibility of the candidate explanation; (ii) that judgements of explanatory value were predicted by the degree of statistical relevance of the candidate explanation for the explanandum.; and (iii) that judgements of explanatory power were sensitive to the framing of the candidate explanation in causal as opposed to non-causal terms.

In the second study, we examined explanations for singular, token events, where exactly one alternative explanation was provided and no other alternative explanation could be easily produced. Experimental participants were confronted with fictitious scenarios and had to rate the quality of a proposed explanation. In contrast with the first study, these hypotheses had a low level of generality and were explicitly about a particular, token- explanandum. We tested again three hypotheses: (i) that judgments of explanatory value could be dissociated from posterior probabilities or other indicators of rational acceptability; (ii) that judgments of explanatory value were positively associated with causal reasoning and a sense of understanding; (iii) that judgments of explanatory value were positively affected by statistical relevance.

Results from the first study provide evidence that for generic types of explanations involving a complex causal mechanism, the prior credibility of the hypothesis and causal framing jointly raise the perceived explanatory value of the hypothesis. Statistical relevance relations has a negligible impact on explanatory value, where there is an unrestricted number of potential explanations, yielding to causal credibility as the main determinant of explanatory value.

Results from the second study provide evidence that for explanations of single events, judgments of explanatory value are highly sensitive to relations of statistical relevance, and are dissociable from posterior probabilities and other indicators of the rational acceptability of the explanatory hypothesis.

Collectively, these findings provide support to the hypothesis that explanation is a complex structure that taps into distinct types of sources of information in different contexts (Lombrozo 2012). They also call for a reassessment of the rationality of explanatory modes of inference like abductive inference (Lipton 2004). Specifically, our findings indicate that two different kinds of probabilistic cues—the credibility of the explanation and the statistical relevance for the explanandum—contribute to explanatory value, albeit in different circumstances. The level of generality of the explanation (and the explanandum) make a crucial difference: for generic (type) explanations, the prior credibility, but not the statistical relevance boosts explanatory value, whereas for individual (token) explanations, explanatory value co-varies with statistical relevance, but not with prior credibility. This indicates that the probabilistic coherence of explanatory modes of inference is context-specific, and the rationality of abductive reasoning should thus be assessed on a case-by-case basis.

We hope our results will promote “the prospects for a naturalized philosophy of explanation” (Lombrozo 2011, 549), contributing to a theory of explanatory reasoning that is both psychologically accurate and philosophically appealing.

References

Halpern, J., and Pearl, J. (2005). Causes and Explanations: A Structural-Model Approach. Part II: Explanations. British Journal of Philosophy of Science 56: 889–911.

Hitchcock, C. (2008). Probabilistic causation. In Edward N. Zalta (Ed.), The stanford encyclopedia of philosophy (Fall 2008 Edition). <http://plato.stanford.edu/archives/fall2008/entries/causation-probabilistic/>.

Lipton, P. (2004). Inference to the Best Explanation (second edition). London: Routledge.

Lombrozo, T. (2011). The instrumental value of explanations. Philosophy Compass 6: 539–551.

Lombrozo, T. (2012). Explanation and abductive inference. In K. J. Holyoak & R. G. Morrison (eds.): Oxford Handbook of Thinking and Reasoning, 260–276. Oxford, UK: Oxford University.

Salmon, W. (1997). Causality and Explanation: A Reply to Two Critiques. Philosophy of Science, 64: 461–477.

Rachel Cooper. The unluckiness of the disordered

Why is that ringworm, panic disorder, and impacted wisdom teeth count as disorders, but that wrinkles, nervousness, and teething do not? We have a strong intuition that a condition can only be a disorder if it is somehow unusual or unexpected. However, cashing out this intuition is surprisingly problematic.

STATISTICAL APPROACHES AND THEIR PROBLEMS

The idea that a condition must be statistically infrequent in order to be a disorder (as held by Taylor (1976) and Kendell (1975)) is clearly inadequate, as it cannot accommodate the possibility of a pandemic. A modified statistical approach might try saying that a disease is statistically infrequent most of the time. But it is doubtful whether this will quite do either. Post nuclear war it seems that disorder might come to be the norm for most humans for quite some time.

Boorse's influential account of disorder (1975) employs a variant of the claim that disorders must be statistically infrequent and faces the same challenges. For Boorse, a subsystem dysfunctions when it fails to fulfil whatever functioning is statistically normal for similar organisms, and thus the idea that disorders must be unusual is built into his notion of dysfunction; for example, it is only because Sarah's blood pressure is high compared to that of other middle-aged women that she counts as having a disorder. This approach again faces difficulties dealing with pandemics; as it seems plausible that all the organisms in my reference class might suffer a dysfunction at the same time.

DEVELOPING A MODAL APPROACH

If we give up on the idea that statistics can inform the distinction between what is normal and what is disorder we might try adopting a modal approach. Cooper (2005), for example, claims that someone can only be said to suffer from a disorder if there are a good number of suitably nearby possible worlds in which they, or their counterparts, are better off. Modal approaches deal better than statistical approaches with the possibility of pandemic: if I have bird flu in a bird flu pandemic my state is statistically normal, but there are still possible worlds where I am better off.

In considering whether there are possible worlds in which I am better off, the worlds to be considered need to be those in which there are humans of something like the actual biological "design". We would ignore far distant worlds in which people live forever, or in which human anatomy has been re-jigged to make giving birth painless. Rather we should focus on worlds in which there are humans designed like us and ask whether someone is badly off compared to them.

Here I go beyond Cooper (2005) in developing the details of the modal constraints under which someone can be said to have a disorder. To deal with the problems caused by genetic essentialism (especially in the face of genetic disorders) I argue that a counterpart approach should be adopting in judging whether I might have been better off.

It would be tempting to claim that someone can only have a disorder if they are better off in most nearby possible worlds, but I argue that this prima facie plausible claim must be rejected for the following reason: There are some creatures whose statistically usual state in the actual world, presumably throughout evolutionary history, is to be diseased (80% of female rabbits develop uterine cancer). In such cases, the evolutionary history of a species becomes tied up with it generally having some disease. In such cases there will then be no nearby possible worlds in which creatures of that "species-design" exist but where the statistically usual state is to be healthy. In the case of rabbits, for example, it seems likely that the high rates of cancer are caused by the hormone levels associated with high fertility. In so far as "breeding like a rabbit" is part of the "species design" for a rabbit, there will thus be no rabbits in nearby possible worlds that are not disposed to develop cancer. I conclude that a rabbit with cancer counts as disordered not because most of its counterparts in nearby possible worlds are better off (they aren't), but merely because some are.

How then can we deal with those disorders that occur at the bottom range of a bell-curve distribution (low IQ, very short stature etc)? With reference to historical examples, I suggest that how unlucky someone has to be to be counted as disordered varies with context. To count as disordered someone must be in a state where they could have been better off, and where their condition is counted unlucky enough to elicit pity, and to seem such that it needs rectifying. Whether someone counts as unlucky enough to count as disordered tends to vary as a function of the ease and expense of treatment. When treatment becomes easier and cheaper we tend to count a larger grouping disordered. In the case of low IQ, for example, as there is no treatment, only those at the very bottom of the bell-curve are counted disordered. In contrast, as depression can be treated fairly cheaply and with some success, even those who are slightly more miserable than average get counted as disordered. In this way, the concept of “disorder” functions rather like “poverty”. To count as poor it must be possible that one could be richer, but the amount of money one needs to count as poor is vague and varies with context. Similarly, to count as disordered it must be possible that one could be better off, but the level of biological functioning that will be counted as disordered varies.

References

Boorse, C. (1975) On the distinction between disease and illness, Philosophy and Public Affairs, 5: 49-68.

Cooper, R. (2005) Classifying madness Dordrecht: Springer.

Kendell, R. (1975) The concept of disease and its implications for psychiatry, British Journal of Psychiatry, 127: 305-315.

Taylor, F.K. (1976) The medical model of the disease concept, British Journal of

Psychiatry, 128: 588-594.

Ana-Maria Cretu. What Good is Realism about Natural Kinds?

Natural kinds realism can be understood as a series of views put forward by those scientific realists committed to kinds that ‘latch onto’ the real structure of the world. Natural kinds are believed to be the best explanatory tool in that they explain why theories featuring those kinds prove inductively and predictively successful. This is, in a nutshell, what I take to be the epistemological argument for natural kinds. A main proponent of this argument is Richard Boyd whose account has become the received view of realism about kinds. In a series of papers ([1991], [1999a] [1999b]) Boyd articulates a realist account of natural kinds: homeostatic property cluster kinds (HPCK). Natural kinds are, on Boyd’s view, necessary to establish the reliability of successful epistemic practices. The idea behind the epistemological argument is the following: we are to some extent justified in giving a posteriori definitions of natural kinds in certain ways that reflect the actual causal structure of the world because we cannot make projectible generalizations otherwise (Boyd [1991], p.138).

Boyd’s HPCK account is designed to explain how kinds used in successful epistemic practices `latch onto’ natural divisions in nature. His main motivation for defending such an account is that of defeating skepticism about the success of science, which is reminiscent of the Lockean nominalist tradition. According to this tradition, “we must classify substances according to arbitrary nominal essences instead of according to microstructural real essences” (Boyd, [1991], p. 131). This is because we cannot know the ‘real essences’ of kinds. Hence, things are classified in virtue of some arbitrary nominal essences. But kinds that are the result of such arbitrary classifications cannot support successful inductive generalizations and make knowledge of such kinds in general seem impossible. The Lockean tradition of kinds thus gives rise to a “tension between empiricist nominalism and the task of accounting for induction” (Boyd, [1991], p. 130). Hence, this tradition is largely responsible for opening the doors to skepticism about the ability of science to use kinds to ground epistemic practices. Boyd’s aim is to revoke such skepticism by showing that “in induction and explanation we must refer to kinds whose definitions are specified a posteriori, in deference to nature, rather than nominally” (Boyd, [1991], p. 131).

In this talk I argue that a realist account of natural kinds à la Boyd is neither necessary nor sufficient to explain success in science. In analyzing Boyd’s account I distinguish between the constitutive factors of HPCK and the individuation conditions of HPCK: these distinctions are crucial for understanding Boyd’s epistemological argument for kinds; they also constitute the basis for the subsequent objections that are aimed at establishing whether or not Boyd’s HPCK account is a genuinely realist account of natural kinds. I argue that Boyd's HPCK account is neither necessary nor sufficient for grounding epistemic practices in science because: i) individuating the constitutive factors of HPCK is more often than not a matter of human decision; ii) the HPCK account falls short of accommodating successful scientific kinds that cannot be described in terms of clusters of properties and underlying homeostatic mechanisms; and iii) the HPCK account includes as kinds things that by the lights of our present science, failed to latch onto the causal structure of the world (for related arguments see Ereshevsky&Reydon [2015], Slater [2014], Khalidi [2013]). Failing to deliver on their epistemic potential, the commitment to HPCK proves not to be the best available tool in the scientific realists' toolbox. I conclude that the epistemological argument should not be made dependent on natural kinds carving nature's joints in some strong realist sense. We can still think of natural kinds as explaining why theories featuring those kinds prove inductively and predictively successful whilst not having just one account of natural kinds. There is not one notion of 'natural kind' that best serves science; in fact the notion of 'natural kind' changes and matures with scientific progress.

Having shown that Boyd's realist account of natural kinds is neither necessary nor sufficient for grounding epistemic practices in science, I conclude that a commitment to natural kinds in some strong realist sense is not necessary to establish the reliability of successful epistemic practices. Instead, we should opt for a less ontologically inflationary account. Taking the cue from Quine's discussion in "Natural Kinds" [1969] I aim to rehabilitate the view that there is no account of natural kinds that spans across all sciences. Instead, I want to suggest, following Quine, that there is a sense in which no unique account of natural kinds is suited to account for the epistemic endeavors of all science. However, there is a sense in which a particular notion of kinds pervades all science. This notion, which is (at least implicitly) deployed by each branch of science in its respective epistemic practices, is the notion of scientifically entrenched kinds. It's worth noting that the notion of 'scientifically entrenched kinds' should be understood as a methodological place-holder for all types of kinds that are useful for the epistemic endeavors of the natural and social sciences. The account I propose is a type of pluralism about accounts of kinds that serve some important epistemic role in science, which accommodates the strengths of the HPCK account, whilst not sliding into Dupre's 'promiscuous realism'.

Erik Curiel. If Metrical Structure Were Not Dynamical, Counterfactuals in General Relativity Would Be Easy

There are important problems with modality in general, and with the understanding of counterfactuals in particular, peculiar to general relativity as a physical theory, arising from the dynamical nature of spatiotemporal structure in the theory. To draw these problems out, consider the following intrepretive principle for general relativity: "in any spacetime, any smooth curve can be reparametrized so as to be a null geodesic iff it could be the trajectory of a light ray." Now, this principle is relatively straightforward to understand when we are considering the possible paths of light rays in vacuo, but how are we to understand the modal force of the claim when matter is present? Surely we want to talk as well about the physical significance of the null cones even at those places. In order to do so, and in order to formulate the analogue of the principle for those spacetime regions (in order to give a physical interpretation to the null cones at those points), we must say something along the following lines: the null geodesics where matter is present are those paths light rays would follow if the matter there were removed. But on its face, that modal statement makes no sense in the context of general relativity, because however we make sense of the idea of "removing matter" from a spacetime region, the metric will eo ipso be different in that region from what it was, and it will generically be the case that the new metric in that region will not agree with the original metric on what it counts as null vectors, much less on what it counts as null geodesics, among many other differences. The distribution of matter in a region of spacetime in large part informs the metrical structure there, so what sense can be made, in the context of the theory, in asking what the metrical structure *would* be if the matter actually there *were not* there?

The problem is made more acute by the fact that metrical curvature is *only* in part informed by the distribution of matter: the Weyl curvature at a point, exactly that part of the curvature encoding conformal information, such as what counts as a null vector, is independent of the value of the stress-energy tensor at that point---the value of the Weyl tensor, point by point, is not constrained by the presence or absence of matter. In regions without matter, moreover, metrical curvature is governed entirely by the Weyl tensor. Still, the Weyl tensor is subtlely related to the distribution of matter at neighboring points, when there is such matter, in a way that can be made precise by using the Bianchi identity formulated using the so-called Lanczos tensor. Thus, in "removing matter" from a spacetime region, there can be no principled way to determine what the "remaining curvature" will be.

One may decide to keep the Weyl tensor the same. But precisely its relation to stress-energy by way of the Lanczos tensor means that this is not an unproblematic way to proceed, and is likely even incoherent or inconsistent. The root of all these problems lies in the fact that there is not a unique vacuum solution to the Einstein field equation.

To make the problem more precise, consider the attempt to take the limit of Schwarzschild spacetime as the central mass goes to 0, because one is interested in the counterfactual question "what would Schwarzschild spacetime look like if its mass were made to vanish?" In Schwarzschild coordinates, using the inverse-third root of the Schwarzschild mass rather than the mass itself, the metric takes a form that clearly has no well defined limit as the mass-parameter goes to zero. Applying one coordinate transformation, the metric takes a form in which the limit does exist and yields a flat solution discovered by Kasner. If instead of that coordinate transformation we apply a different one to the original Schwarzschild form, then the resulting form also has a well defined limit, which is the Minkowski metric. Thus, the limits in the different coordinates yield different metrics, with no natural or preferred way to say which is the "correct" limit. (Using the geometrical machinery developed by Geroch for taking limits of spacetimes in an invariant way, this argument can be made precise and rigorous without reliance on coordinate systems.)

Compare the situation in Newtonian gravitational theory. It makes perfect sense in Newtonian theory to reason counterfactually about the behavior of a given kind of system in the presence or absence of any other kind of system, since that presence or absence won't affect the kinematical structure of Newtonian spacetime. There is, for example, no problem in principle in computing the counterfactual change in gravitational forces in a region induced by any counterfactual changes in the distribution of matter anywhere in the spacetime. But one just cannot do that in general relativity, unless one spells out what the new metrical structure will be in advance when one tries to reason counterfactually about what would happen if one were to "change the distribution of matter in a region of spacetime". But there is no canonical or natural way of spelling out the metrical structure in advance.

The problem I expose in this paper is severe: many influential philosophical approaches to many fundamental problems and issues in the philosophy of science---the nature of scientific laws, of theory-confirmation, of causation, et al.--- rely, in ineliminable ways, on subjunctive conditionals for their formulation and application. Physicists certainly rely on such propositions in theoretical and experimental practice to propose and perform tests of general relativity. What reason do we have to believe that we understand what is happening in such cases in the context of general relativity, much less to have confidence in any conclusions drawn?

Indeed, I think the situation is even worse than the preceding remarks suggest. Because the problem arises solely from the dynamical nature of spacetime geometry in general relativity, what I say here is wholly independent of one's favorite account of counterfactuals---it depends only on the theoretical resources general relativity provides to model such situations and pose such propositions, no matter what ancillary tools or frameworks one uses to interpret and understand them.

Haixin Dang. Theory choice during conceptual change: The case of WH Bragg and X-rays

In this paper, I discuss an important methodological problem in science: How do we choose between theories during periods of conceptual change? Philosophers of science have spent a lot of time analyzing theory choice under retrospective, idealized circumstances. For example, the problem is often framed with two complete theories which offer competing interpretations for a complete set of experimental results. The existing literature on theory choice has largely revolved around the problem of underdetermination of theory by evidence. Here I will offer a different perspective by focusing on conceptual change. When scientists hit the boundaries of the explanatory power of the present physical theories, as during the time of major conceptual change, how should we interpret experiments with incomplete and competing theories?

The view I will come to defend throughout this paper is what I call “pragmatic instrumentalism.” This is the view that during times of theory change, it is rational for scientists to become instrumentalists towards their theories. I will reason through a historical case to fully illustrate this position. While theory choice is an important topic for philosophers of science, I argue that philosophers need to examine how theory choice works in practice, that is during actual periods of conceptual change, to really understand the rationality of science. The case study I use in this paper is one of the major conceptual changes of the 20th century: the wave-particle duality of light. Understanding of this phenomenon came slowly and required major revisions in the frameworks of physics. William Henry Bragg played a major role in this early history. I argue that Bragg adopted a “pragmatic instrumentalist” position towards the competing wave and particle theories of light before the emergence of quantum theory in the 1920s.

In the first two decades of the 1900s, the dual nature of light was being debated and Bragg most keenly felt the difficulties in reconciling the two theories. He had spent most of his career working on X-rays and, for many years in the early decades, was one of the foremost defenders of a corpuscular interpretation: the neutral pair hypothesis. Bragg, however, will always be most well known for providing the strongest proof for the wave nature of X-rays. For his work in the analysis of crystal structure by X-ray diffraction, Bragg received the Nobel Prize in Physics in 1915, an honor that was shared with his son William Lawrence Bragg.

Since the discovery of X-rays in 1896, the strange behaviors of these new rays have perplexed physicists. In the early decades, the relationship between light and X-rays was unclear and whether X-rays were waves or particles was up for debate. When Einstein wrote his 1905 paper on the photoelectric effect, he was concerned with ultraviolet light, not X-rays, and most of the X-ray researchers, especially outside of Germany, were skeptical of the light quantum; many did not believe light and X-rays to be the same phenomena. After Laue devised his 1912 experiment demonstrating X-ray diffraction by crystals, X-rays were then understood as a kind of light. But the other perplexing properties of X-rays remained and caused physicists to reconsider the nature of all electromagnetic waves; these problems were not fully explained until the 1920s. Bragg played an important role in this early history. His debate with Charles Glover Barkla in 1907-1908 over the nature of X-rays was the first wave-particle controversy of the century. Bragg’s insights into the behavior of X-rays were very prescient of wave-particle duality. While Bragg had no contact with Einstein, he still became one of the first advocates of a “quasi wave-particle” theory.

I will argue in this paper, that throughout this time of controversy, Bragg viewed his theory as a working model, that is, he held an instrumentalist view of the corpuscular theory. I will argue that this way of understanding Bragg’s commitment to the corpuscular theory explains both how he was able to defend his theory prior to 1912 and also explains why he continued to hold the view after 1912. By looking through Bragg’s published papers, as well as, his private letters held at the Royal Institution of Great Britain and the University of Cambridge, I show that at the very conception of his neutral pair hypothesis, Bragg held a physical model that contained both wave and particle elements. I argue that over the course of his debate with Barkla and as new X-ray phenomena surfaced over 1908-1911, Bragg had changed his view to advocate a working model, essentially giving up claims to the earlier physical picture. It is in this sense that I mean pragmatic instrumentalist: a working model that can be exploited for constructing hypotheses and experiments, but does not claim to be physically true. I also argue that Bragg was surprisingly consistent in his view of the instrumental importance of the particle theory even after Laue's experiments. He continued to believe that the corpuscular model captured something that the accepted wave model was missing.

This episode in the history of science illustrates interesting philosophical results regarding theory and experiments. Throughout the controversy, Bragg found himself hitting the boundaries of the explanatory power of the present physical theories. I will argue that Bragg's position, pragmatic instrumentalism, is a rational one in face of scientific controversy. Bragg's emphasis on pragmatic concerns— especially how fruitful the theory is to the development of future research—is a rational criteria to hold. While his contemporaries harshly criticized Bragg for holding on to his corpuscular theory, I argue that ultimately Bragg was being a "good" experimenter in maintaining the conviction in his results and holding his theory and the wave theory to a higher explanatory standard.

Finally, in the last section of this paper, I argue that pragmatic instrumentalism is not only useful in analyzing the history of science, but also a useful philosophical position. It allows us to understand the rationality of theory choice during times of conceptual change.

Radin Dardashti. No Alternatives for What? Non-empirical Evidence in the Case of String Theory

In a recent paper Dawid, Hartmann and Sprenger have shown within a Bayesian framework that the observation that there is no alternative theory to one's theory, at a given time and despite considerable effort, confirms the theory. This so-called No Alternatives Argument (NAA) is crucial in cases where empirical evidence is missing, as in String Theory. Unlike common theory confirmation the confirming evidence in this case is called non-empirical, since it is not a deductive or inductive consequence of the theory that there are no alternatives. The main focus of this paper is how one can obtain such non-empirical evidence in an objective, unbiased way. The conclusion is that in cases where the NAA is most needed (i.e. in theories where empirical evidence is missing), it is usually not yet applicable, while in cases where one does have enough non-empirical support, empirical evidence can be given too (as in the case of the Higgs mechanism), and so the NAA is not needed. The paper is divided into three parts, which I will discuss now.

1. What is the precise definition of non-empirical evidence in the NAA?

In the first part we critically analyse the definition of non-empirical evidence in the NAA and argue that its formulation in Dawid et al. (2015) is inadequate for the purposes of theories of quantum gravity, i.e. for those cases where it is most needed. We offer an extension and a problem-relative reformulation of the definition of non-empirical evidence, which allows for an application of the NAA to the relevant theories. Any NAA is then always relative to the specific set of problems P the theory is meant to solve. There remain two open questions: First, how do we individuate theories? And second, what is the specific problem set?

2. How to individuate theories?

The first problem arises from the need to individuate theories. It is obviously crucial for the No Alternatives Argument that it be possible to claim that there are no alternatives to one's theory. This, however, implies the possibility to individuate theories, since only if I can count theories, can I claim that the number of theories solving a problem is one. After arguing why the answer offered by Dawid et al. (2015) is not satisfactory I propose an alternative criterion of theory individuation, which for the purposes of the NAA offers a pragmatic solution to the problem. This criterion offers a problem-relative individuation, since different problem sets can lead to different individuations. I will consider several examples to illustrate the applicability of this criterion.

3. What is the right problem set?

The second problem is due to the problem-relative statement of the non-empirical evidence in NAA. If one says one has no alternative, one always needs to specify with respect to what problem set there is no alternative. E.g. in the case of String Theory the statement is that it “is the only viable option for constructing a unified theory of elementary particle interactions and gravity”. But who determines what the relevant problem in need of a solution is? The determination of this set of problems is a priori highly non-trivial, especially in the cases where the NAA is most crucially needed, where the determination of the problem set can be dependent on the research program within which the scientist works. While we offer a pragmatic solution to the theory individuation problem, the problem-determination problem remains and leaves us with two possible interpretations of the NAA result in this light:

The first possible interpretation follows from the fact that if there is no way to justify the problem set independently, any scientist may regard her own favorite set of problems. This seems especially adequate in the context of theories of quantum gravity, where each research community has their own set of problems and favorite methods by which they aim at solving them. This has the following more general and rather undesirable consequence: within each research project one can find a unique problem set such that (according to the criterion of theory individuation) there will be no alternatives to that theory. What is the meaning of the confirmatory result of Dawid et al. in this light? The argument does not trivialise completely, since scientists work on the specific theories they are working on because they consider the theory they use as most appropriate considering the set of problems they wish to address. If there were many alternatives able to address the same problem set, their trust in their specific approach may decrease. The confirmatory result that follows from the NAA should then be understood as a justification for the scientists to work on the theory they use given their specific problem set. The confirmatory result should then not be understood as a result confirming the theory per se but as a result which accounts for the scientific practice.

The more interesting conclusion would be that there is a preferred problem set. In this case the NAA by itself may provide theory confirmation. Consider, however, the unification of all fundamental forces. Whether or not this should be considered as a problem in need of an explanation is non-trivial. So these claims go beyond the empirically justified problems. So if they are not empirically justified, one can only evaluate them by considering the appropriateness of the assumptions within the bigger research program. For instance, if unification has been the right guide in the development of theories in the past then they may be in the future as well.However, I argue that this kind of meta-inductive support is not available for theories of quantum gravity. This may lead to the unfortunate consequence that the NAA in cases it is most needed, it usually will not yet be applicable, while in cases where one does have enough non-empirical support, empirical evidence can be given too, and so the NAA is not needed.

Hugh Desmond. Natural Selection: Convergence and Causality

Until recently it was relatively uncontroversial to say natural selection is one of the causes driving evolution. In fact, in biology textbooks natural selection is often represented as some kind of Newtonian force, with magnitude and direction. However, this picture is complicated when one takes the statistical nature of selection into consideration. Evolution by natural selection is constituted by individual births and deaths, and strong arguments have been developed that selection cannot be some causal propensity over and above individual-level processes. Following a number of articles by Walsh, Ariew and Matthen, there is now a counterposition that natural selection is a mere book-keeping of the genuinely causal interactions that take place between individual organisms. It is not a cause, let alone a

Newtonian force (Matthen and Ariew 2009; Walsh, Lewens and Ariew 2002; Walsh 2007).

In the extensive literature that has ensued, the statisticalist approach has mainly been used to argue for a deflationary position: “fitness and natural selection have no reality except as accumulations of more fundamental events” (Matthen and Ariew 2002, 82). In this paper I will to investigate the underexplored possibility of a non-deflationary statisticalist analysis of selection. This adopts the statisticalist, bottom-up analysis of population change, but tries to reconcile it with certain causalist intuitions. The inspiration for this is that, while statisticalist considerations may preclude certain naïve ways of understanding the causal nature of selection, causalist intuitions cannot be entirely wrong either. At the very least, it cannot be denied that most of biological practise is not threatened by these considerations. While it may be metaphysically inacccurate, it is often empirically accurate to model selection as a causal force (for example in cases of stabilizing selection, where component pressures cancel out). This suggests that causalist intuitions must be legitimate in some way.

My approach in this paper will be to use the notion of equilibrium as a way of understanding how the causal nature of selection can be real, thus grounding causalist intuitions. Equilibrium is a central concept in modeling the behavior of complex systems. In particular, stable equilibria are empirically important because they act as attractors and allow for a long-term prediction of the behavior of the system, even though the behavior in the middle-term may be chaotic and too complex to calculate. However, they are also philosophically important as they can allow a well-defined direction to be assigned to a complex process. Thus a concept of directionality can be formulated that is grounded in a statistics of individual-level dynamics and that allows us to understand why natural selection can be legitimately called causal.

To establish such a framework, I will need to do three things. The first task will be to lay the ground by disentangling some different notions of causality at play, in particular process and difference-making causality. Each highlights a different aspect of natural selection and confusion results if these are not kept separate. In this paper I will focus on difference-making causality alone, mainly because this notion has been more controversial. Difference-making is, broadly, counterfactual dependence. The statisticalist arguments have endeavored to show that, even if natural selection were not present, evolutionary change would occur.

One argument has been that natural selection is established only retroactively, by a statistical regression on actually occurred births and deaths (where selection is the correlation between traits and births). There is no description-independent way of establishing fitness or natural selection (and this is related to the reference class problem). Another argument has concerned the inseparability of natural selection from the causal processes affecting the behavior of organisms. The probabilities that characterize the possible outcomes by natural selection are only a measure of our ignorance of the individual-level processes determining the births and deaths. They do not correspond to any putative ‘causal propensity’ that could be used to ground natural selection.

The second task will be to formulate the condition of equilibrium, and to show how, if it is accepted, it can resolve certain key issues regarding difference-making causality. For this I will use an extension of the Price equation to the multigenerational case. The Price equation gives an exact relationship between the phenotype distribution of different generations, and I will show how this equation can be simplified considerably under assumption that an equilibrium is reached after a certain number of generations. This assumption then allows one to uniquely define a direction of an evolutionary process: the tendency towards equilibrium.

This is important because it allows one to argue that the probabilities defining fitness are not purely description-dependent. Neither is natural selection merely a measure of subjective uncertainty; rather, it reveals an objective feature of certain evolutionary processes, namely the presence of stable equilibrium. Natural selection is causal in the difference-making sense: if it were not present, an evolution towards stable equilibrium would not be observed.

Finally I will need to argue why the equilibrium condition is a plausible assumption. To this end, I will show that given evolutionary change, either a stable equilibrium is reached, or if it is not, then the concept of fitness is not meaningful. I discuss certain results from Markov process literature, where the conditions for equilibrium are established (Doeblin's theorem). From this it can be seen that the notion of equilibrium is intertwined with natural selection, and that this is a natural way to reconcile both statisticalist and causalist approaches.

Neil Dewar. Symmetry, differences, and naturalism

This paper is concerned with the claim that the presence of symmetries in a physical theory impose prima facie constraints on how that theory is best interpreted: in particular, with the claim that models of a theory related by a symmetry should be interpreted as representing the same state of affairs. It contends that a formal characterisation of symmetries is sufficient to ground this claim.

The first half of the paper considers this as a question of theory interpretation, i.e., as a matter of drawing out a theory's ontological commitments. First, the paper outlines how symmetries may be formally specified, as fibre-preserving transformations which map solutions of the theory to other solutions. It then argues that such transformations may be understood as codifying the notion of differences between models, and uses this to argue that symmetry transformations correspond (in a precise sense) to formal differences without a physical correlate: that is, that symmetry transformations are precisely those differences which, although "visible" to the theory's formalism, are irrelevant to the theory's dynamics.

The remainder of the paper concerns how the above can be used to draw substantive metaphysical conclusions: that is, how to pass from a maxim for interpreting theories to a justified belief about what the world is like. I consider two strategies for accomplishing this. The first (the "epistemic strategy") turns on the claim that structures which do not figure in physical laws in the right way cannot be detected by beings whose epistemic processes are governed by those laws. Making this idea precise turns out to require a careful analysis of the relationship between knowledge and action, and of the role of physical law in individuating those actions.

The second strategy (the "nomological strategy") advances a novel conception of the relationship between laws and ontology. Rather than taking laws as being statements constraining the behaviour of an antecedently given ontology, we may take ontology be a codification of the structures articulated in the laws. (I take this to be one way of spelling out the idea of "ontic structural realism".) By doing so, we are bound to posit ontology only insofar as it is needed to encode such nomic structure; hence, having interpreted a theory as expressing a given set of laws, we are enjoined to rule out any structures which are irrelevant to that project of encoding.

I conclude with some remarks about possible future directions for this research, and how the ideas expressed here related to broader issues in philosophy of science (concerning, in particular, naturalism, formal approaches to theoretical equivalence, and structural realism).

Joe Dewhurst. Natural kinds and folk kinds in the psychological sciences

This paper will examine the role that natural kinds play in psychology and cognitive science, and ask whether folk psychological kinds are capable of fulfilling this role. I will first specify what I mean by natural kinds and folk psychological kinds, and then argue that the latter are not suitable for the job required of natural kinds in the psychological sciences. Whilst folk psychological kinds constitute what Hacking calls "human kinds", this is insufficient to qualify them for full natural kind status, even in the limited capacity outlined in this paper. Furthermore, the use of folk psychological kinds threatens to systematically undermine both theoretical and experimental work in psychology and cognitive science.

For this reason, I will conclude that a concerted effort is required in order to develop new conceptual categories that more accurately reflect our understanding of the human cognitive system.

Natural kinds terms play a central role in scientific discourse and practice, regardless of whether or not they are referred to as such. By this I simply mean that the projectable predicates required for inductive inference resemble what we typically think of as natural kinds (cf. Quine 1970). This fact alone does not entail any stronger claims about the ontological or metaphysical status of natural kinds. It is also important to acknowledge the pragmatic (or perhaps sociological) importance of natural kind terms (Wikforss 2010, Brigandt 2011, and Khalidi 2013 come to similar conclusions), even if one were not interested in the broader philosophical debate.

It is typically the case that the projectable predicates deployed by a science will, in the first instance, follow the example set by intuitive folk taxonomies (Gopnik & Schwitzgebel 1998: 78-9). In physics and chemistry we began with the observable properties of objects, in biology we began with obvious environmental and physiological groupings, and in psychology and cognitive science we typically begin with folk psychological taxonomies. A key difference here is that whilst physics, chemistry, and biology have all at least partially transcended their folk taxonomical beginnings, in the psychological sciences we are by and large still stuck with folk psychology. We must ask, therefore, whether the folk psychological taxonomy is fit for purpose.

Whilst there is no general agreement as to which account of natural kinds is correct, it is at least broadly acknowledged that to be fit for purpose in the biological and psychological sciences, an account of natural kinds should allow for a degree of flexibility in membership conditions. Either we find such an account, or we must conclude that the kinds of biology and psychology are not natural kinds. A promising candidate for such an account is some version of the homeostatic property cluster theory, which claims that (at least some) natural kinds consist of regularly co-occurring clusters of properties along with a homeostatic mechanism that explains the co-occurrence of those properties (see e.g. Kornblith 1993: 35, Boyd 1999, Magnus 2012). Accounts of this kind are fairly liberal, and for the purposes of this paper I will take them as a yardstick against which to measure the success of folk psychological kinds. If they fail here, then they are unlikely under any more stringent account of natural kinds.

There are two reasons to think that folk psychological kinds might not be natural kinds. The first has to do with the extent to which folk psychological explanation and discourse varies across cultures and languages. Given that different cultures draw on different taxonomies when attributing mental states (see e.g. Lillard 1998, Turner 2012), it seems that we cannot simply read off a 'correct' taxonomy that will correspond to the natural kinds of psychological science. Of course it might be the case that genuine psychological kinds will correspond to some folk psychological kinds, but, prior to experimentation, there is no way of knowing which these will be. We certainly cannot assume that the folk psychological kinds of our own culture or language will correspond precisely to the kinds of a finished psychological science.

The second reason for thinking that folk psychological kinds are not natural kinds will apply even if one was able to uncover some cultural universals that were not vulnerable to my first argument. By and large, folk psychological kinds are not suitable for fine-grained scientific enquiry. Consider the archetypal folk psychological kinds, belief and desire. Whilst they are prevalent in philosophical thought experiments, these terms rarely feature in scientific psychology. When they do appear, they are used to refer to a far more disparate set of concepts than the folk kinds encompass (see e.g. Krueger & Grafman 2013). This means that folk psychological kinds are disjunctive in a way that is ruled out by most contemporary accounts of natural kinds (see e.g. Khalidi 2013: 89-92). Without further refinement, folk psychological kinds are not suitable for the role required of natural kinds in the psychological sciences (i.e. projectability across different domains).

Given that folk psychological kinds appear not be natural kinds, what kind of a thing are they? They certainly appear to be projectable in at least some non-scientific context, such as when they are used to predict the coarse-grained behaviour of conspecifics. It is explanatory power in this sort of context that defenders of folk psychological kinds tend to appeal to. However, it is also precisely this sort of context that introduces the problems raised by Hacking with regard to what he calls "human kinds" (1995). Folk psychological kinds are only projectable in social contexts, where they are dependant upon the looping effects described by Hacking and more recently explicated by Zawidzki (2013) as "mindshaping". That is to say, folk psychological kinds only have explanatory power when the very act of using them enforces their own validity by shaping the way in which we behave and think. They lose this explanatory power as soon as we descend below explanation in the social domain, and as such are ill suited for the role required of natural kinds in any more general account of psychology and cognitive science. We must therefore look elsewhere for a psychological taxonomy that is fit for purpose.

References

Boyd, R. 1999. "Homeostasis, species, and higher taxa." In Wilson (ed.), Species: New Interdisciplinary Essays. Cambridge, MA: MIT Press.

Brigandt, I. 2011. "Natural Kinds and Concepts: A Pragmatist and Methodologically Naturalist Account." In Knowles & Reidenfelt (eds.), Pragmatism, Science and Naturalism. Peter Lang.

Gopnik, A. & Schwitzgebel, E. 1998. "Whose Concepts Are They, Anyway? The Role of Philosophical Intuition in Empirical Psychology." In De Paul & Ramsey (eds.), Rethinking Intuition. Lanham: Rowman & Littlefield.

Hacking, I. 1995. "The Looping Effects of Human Kinds." In Sperber, Premack, & Premack (eds.), Causal Cognition, an Interdisciplinary Approach. Oxford, UK: OUP.

Khalidi, M. A. 2013. Natural Categories and Human Kinds. Cambridge, UK: CUP.

Kornblith, H. 1993. Inductive Inference and Its Natural Ground. Cambridge, MA: MIT Press.

Krueger, F. & Grafman, J. (eds.) 2013. The Neural Basis of Human Belief Systems. Hove, UK: Psychology Press.

Lillard, A. 1998. "Ethnopsychologies." Psychological Bulletin, 123/1: 3-32.

Magnus, P.D. 2012. Scientific Enquiry and Natural Kinds. Basingstoke, Hampshire: Palgrave MacMillan.

Quine, W.V.O. 1970. "Natural Kinds." In Rescher et al (eds.), Essays in Honor of Karl G. Hempel. Dordrecht: D. Reidel.

Turner, R. 2012. "The need for a systematic ethnopsychology." Anthropological Theory, 12/1: 29-42.

Wikforss, A. 2010. "Are Natural Kind Terms Special?" In Beebee & Sabberton-Leary (eds.), The semantics and metaphysics of natural kinds. London: Routledge.

Zawidski, T. 2013. Mindshaping. Cambridge, MA: MIT Press.

Janette Dinishak. Autism, Aspect-Perception, and Deficit Explanations of Human Differences

In this talk I argue that there are significant problems with how philosophers and other theorists approach the study of autism. My main concern is that many recent approaches primarily understand autism in terms of deficits. Roughly, this involves conceptualizing autism as the lack or absence of some feature, trait, capacity, etc. and then characterizing this lack or absence as a deficit in the feature, trait, capacity, etc. in question (i.e., as the lack of absence of some feature, trait, capacity, etc. that one ought to have). However, deficit-based approaches to understanding phenomena have historically proven dangerous and problematic in a variety of cases. I begin by describing examples from philosophy and science that illustrate some ways in which deficit-based approaches can be problematic. Then I articulate and assess the deficit treatment of autism. I do not claim that deficit-based approaches are never appropriate; nor do I claim that they will definitely prove harmful in the case of autism. But the dangers of deficit-based approaches to understanding autism are significant enough to warrant our proceeding extremely carefully.

Perhaps the most well-known and influential instance of a deficit treatment of autism is the "theory of mind" account. On this view autists have a specific cognitive deficit: a lack or delay in the development of the "theory of mind" module. It is hypothesized that a theory of mind deficit explains autistic individuals' social and communicative difficulties since this module is supposed to account for typical individuals' ability to attribute mental states (e.g., intentions, beliefs, desires) to oneself and to others, an ability thought to be integral to explaining and predicting behavior. My focus is another application of a deficit-based approach to understanding autism—the recent appeal by philosophers of mind (Overgaard 2006; Stawarska 2010; Proudfoot 2013) to the notion of "aspect-blindness" to explain autists' difficulties with social interaction.

The idea of "seeing aspects" originates with and was developed by the philosopher Ludwig Wittgenstein (1953/2009b). One kind of aspect is what I call "psychological aspects". Ordinarily, in many cases, one can see emotion in the faces and bodies of other people: in their facial expressions, tones of voice, gesture, posture, and gait. One can see a glance as an expression of shyness, or hear a plea as hesitant and so forth. Wittgenstein also introduced the idea of "aspect-blindness." A person blind to psychological aspects would be unable to see a person's narrowed eyes and downward turned mouth as an expression of anger or to hear a voice as joyful, for example. They are thereby said to be "blind" to the angry facial expression or to the joyfulness in the voice. While more attention to the perceptual dimensions of autism is a welcome development in philosophical explorations of the condition, I argue that this rendering of the relationship between autism and aspects is problematic. Philosophers should broaden their frame for understanding autism beyond aspect-blindness to include aspect-perception. I discuss two closely related reasons for this recommendation.

First, the science and philosophy of autism is young. Our understanding of the nature of sensory-perceptual differences and atypical social cognition in autism is tentative at best. Sensory-perceptual differences in autistic individuals are proving to be challenging to describe, measure, and relate to autists' social and communicative atypicalities. To date, only a few perceptual phenomena have been systematically investigated. In addition, determining the profile of autists' cognitive and perceptual strengths and weaknesses has been limited by small sample sizes and an absence of comprehensive behavioral phenotype information in autists (Charman et al., 2011).

Second, even if it turns out that empirical findings on autistic perception decisively support the attribution of some forms of aspect-blindness to autists, a frame for capturing autistic experience that only makes use of the notion of aspect-blindness is too narrow. Theorists of autism should broaden out to the notion of aspect-perception, which, unlike aspect-blindness, makes room for conceptualizing autists as having points of view on the world that are not simply a matter of missing things. A focus on aspect-perception provides us richly descriptive tools to understand autistic experience in terms of sensory-perceptual differences rather than merely as a form of deprivation. Such an emphasis could help theorists sharpen and assess the idea (put forth by supporters of the neurodiversity movement) that being autistic involves unusual but not deficient ways of being in, experiencing, and knowing the world. Autists may, for instance, attend to different perceptual features and mobilize different concepts in their perceptual experience of people, objects, and environments, enabling them to perceive aspects "neurotypicals" do not perceive. Further, autistic autobiographies teach us that it distorts the phenomenology of autistic experience to characterize it predominantly in the language of deprivation, suffering, and severity. By focusing exclusively on connections between autism and aspect-blindness, philosophers prematurely close off a whole host of possibilities for understanding autism that may be available to us if we approach the study of autism with a broader framework, one that includes conceptualizing autists as engaging in alternate forms of aspect-perception.

One upshot of this talk is that philosophy must be engaged with and informed by an understanding of autistic experience. Autists have perspectives on the world that are not just a matter of missing things. Another upshot is that deficit views have historically proven dangerous and problematic in a variety of cases and the dangers of deficit-based approaches to understanding autism are significant enough to warrant our proceeding extremely carefully. In the case of autism and perhaps more widely, deficit views may be socially harmful, but they may also impede progress in our understanding of the phenomena themselves. Thus articulating and assessing deficit views such as the "autists are aspect-blind" view is of practical and philosophical importance.

Callum Duguid. Best system accounts and metalaws

While laws are supposed to govern or describe patterns in the non-nomic facts that obtain at a world, metalaws are supposed to govern or describe patterns found in the laws of a world. The practice of science suggests several plausible candidates for this latter role, for example: that the laws must Lorentz-invariant, that they are invariant under spatial displacement and that they do not vary across time. Marc Lange has challenged advocates of the Best System Account of lawhood (BSA) to find a way in which they can accommodate these metalaws in their account. In short: if the laws are just the generalisations found in the deductive system that strikes the best balance between strength and simplicity, what are these metalaws?

There is, in fact, a natural extension of the notion of a BSA-law, and Lange offers it to defenders of the BSA. If the laws are just the generalisations in the best system concerning the world's non-nomic facts, then the metalaws are the generalisations in the best system concerning the world's nomic facts. There are, Lange claims, two issues with this. The first is the problem of counterfactual resilience. Metalaws are supposed to be highly resilient, since scientific practice suggests that when we imagine a world with different laws, we hold the metalaws fixed. But it turns out that Lewis' account of even ordinary counterfactuals has close possible worlds that violate the metalaws, and that doesn't make them very resilient at all! The second problem is that best systems must be formulated in a language whose predicates refer to perfectly natural properties. But the sort of predicates that appear in the metalaws do not plausibly refer to such properties. So a system containing these metalaws will not be ranked as the best.

I am primarily concerned with the second problem. There is a straightforward response to the first: namely that Lange is mistaken in how a defender of the BSA will view closeness of worlds. But the seemingly natural response to the second- widen the notion of natural properties to include those that feature in the metalaws- is fraught with difficulty due to the many different roles that natural properties are supposed to play. Instead, I pursue an alternative approach: abandon the language of natural properties that Lewis appealed to and find a different language to formulate candidate best systems in.

There are some restrictions here, not just any language will do. We need to avoid trivialising the account with Lewis' infamous predicate F that refers to all of the world's truths and has 'for all x, Fx' as its best system. We also need to restrict all competitors for a single best system to a single language. As Cohen and Callender have noted, comparisons of simplicity are relative to a single language so we cannot compare candidate systems formulated in different languages. Two accounts of laws in the literature can be extended so as to cover metalaws.

The first is Cohen and Callender's Better BSA. On this view, the laws are still regularities of the best system, but best systems are language relative. Every language holds its own competition for best system. Those that win give the laws relative to that language, but it simply doesn't make sense to ask what the laws are simpliciter. With an appropriate choice of language, there will be a best system for that language whose laws have the same content as the metalaws we wish to account for. Still, there are two issues with this option worth examining. First, we start to lose track of the connection between the metalaws and the laws on this view. Since both are just laws of systems in different languages, a system that gives us the metalaws may not acknowledge there being any laws, and vice versa. Second, the central issue with the Better BSA is still a concern: if laws are language relative, then concepts that rely on the laws are also language relative. And in cases like evaluation of counterfactuals, that seems mistaken.

The second option is to extend a suggestion by Loewer. Motivated by concerns regarding the suitability of the language Lewis favoured, Loewer suggests that the laws should be assessed in the language of an idealised scientific community. This would allow advocates of the BSA to appeal to a single language that, we can hope, contains the right predicates to formulate metalaws while not containing Lewis' spurious predicate F. One might worry, however, that there may not be one single language that ideal scientists would use, and, should there be multiple ones, choosing one over the others looks unacceptably arbitrary. This is a symptom of a deeper issue, namely that it is not entirely clear how we are supposed to delineate this idealised community, and to what extent the practices of idealised scientists are continuous with those of our own scientific community.

None of the points made above demonstrate that variants of the BSA cannot be made to accommodate to notion of a metalaw; that would require a stronger argument. But they are intended to show that there are difficulties with each option that we can identify and then, hopefully, begin to resolve.

Anna-Maria Asunta Eder. In Defense of a Credence Interpretation of Probability

Williamson argues that an evidential probability (i.e., the probability of a proposition p on an agent s's total evidence) can neither be (adequately) interpreted as the credence of a human agent nor as that of an (epistemic-)ideal agent (2002: 209–11). He concludes that no interpretation in terms of credence is adequate.

If Williamson is right in his criticism of credence interpretations, this has far-reaching consequences for philosophy of science and formal epistemology. It is common among philosophers of science and formal epistemologists to interpret (evidential) probabilities in terms of credences. Remarkable advances in these areas must be withdrawn if such an interpretation is inadequate.

Williamson's criticism of credence interpretations is widely ignored. In this presentation, I intend to make up for this. In the first part, I argue that Williamson's criticism is flawed, and in the second part, I propose a credence interpretation that he does not consider.

1 The Ideal-Agent-Credence Interpretation

Williamson is right in that evidential probabilities cannot be interpreted as the credences of human agents—this is uncontroversial. However, let us consider his argument against interpreting evidential probabilities as credences of ideal agents.

Suppose we are speaking of a specific (abstract) ideal agent si with a credence function Crsi. Then, according to the ideal-agent-credence interpretation, the following holds:

(I): The probability of p on e for s equals r iff it is necessarily the case that [the credence si assigns to p equals r, given e is si's total evidence].

1.1 Against (I)

In the first part of my presentation, I present and criticize what I consider to be the best reconstruction of Williamson's argument against (I). (For the sake of simplicity, I refer to this reconstruction as Williamson's argument.) The argument's starting point is the following quotation:

"[. . .] let a be a logical truth (a proposition expressed by a logically true sentence) such that in this imperfect world it is very probable on our evidence that no one has great credence in a. [. . .] Let b be the hypothesis that no one has great credence in a. By assumption, b is very probable on our evidence [e*]" (Williamson 2002: 209–10; notation adjusted).

In accordance with this quotation and assuming that a is a complex proposition, the first premise runs as follows:

(P1I): The probability of b on e* for s is high (i.e., equal or above a specific appropriate threshold).

The following second premise is true for ideal agents qua being ideal (see Williamson 2002: 210):

(P2I): If p and q are logically equivalent, then it is necessarily the case that the credence the ideal agent si assigns to p equals the credence si assigns to q.

Furthermore, Williamson assumes that it is characteristic for any ideal agent that she does not have great credence in propositions that are of a Moore-paradoxical form. We are led to the following premise:

(P3I): It is necessarily the case that [no matter what si's total evidence is, the credence the ideal agent si assigns to (a&b) is not high].

(P1I), (P2I), and (P3I) together with (I) imply:

(CI): It is necessarily the case that [given that e* equals s's total evidence, the credence the ideal agent si assigns to (a&b) is high and not high at the same time].

(CI) is unacceptable and makes (I) pointless.

1.2 No One?

(P1I) and (P3I) involve b, that is, the proposition that no one has great credence in the logical truth a. It is not at all clear what is meant by 'no one' in the present context. Does it mean the same as 'no human agent' or does it mean the same as 'no human and no ideal agent'? I argue that (P1I) and (P3I) suggest different readings of 'no one' and that, thus, Williamson's argument is not sound.

Even if the exact reading of 'no one' were not that relevant, there is another way out for advocates of credence interpretations. Advocates of credence interpretations might suggest that evidential probabilities should be interpreted in terms of credences of ideal agents who do not have the introspective ability or the respective evidence that includes information about their own epistemic states, thereby denying (P3I). Such a modification strikes me as ad hoc. Instead, one might try to avoid invoking ideal agents altogether.

2 The Ought-Credence Interpretation

In the second part of my presentation, I argue for an interpretation in terms of rational credences that are understood as the credences the agent in question ought to have. (The kind of ought that I have in mind helps us to avoid taking reference to ideal agents. Furthermore, the interpretation is neutral with respect to the debate on whether subjective or objective Bayesianism is adequate; it does not require that there is one probability function that all agents ought to have.) In this vein, what I call the ought-credence interpretation claims:

(O): The probability of p on e for s equals r iff it ought to be the case that [the credence s assigns to p equals r, given e is s's total evidence].

I rephrase Williamson's argument accordingly:

(P1O): The probability of b on e* for s is high.

(P2O): If p and q are logically equivalent, then it ought to be the case that the credence s assigns to p equals the credence s assigns to q.

(P3O): It ought to be the case that [no matter what an agent s's total evidence is, the credence s assigns to (a&b) is not high].

Finally, I show that (P1O) is to be rejected and (O) can be saved. I do so by arguing in favor of the following two assumptions:

(A1): It ought to be the case that the credence s assigns to a is high.

(A2): It ought to be the case that [if s assigns a great credence to a, then s assigns a low credence to b, i.e., the proposition that no one (i.e., no human agent) has great credence in a].

Together with (O), (A1) and (A2) imply the falsity of (P1O). Concluding, advocates of credence interpretations need only deny this premise and endorse (O).

References

[1] Williamson, T. 2002. Knowledge and its Limits. Oxford University Press.

Matthias Egg. Do We Need a Primitive Ontology to Make Quantum Mechanics Empirically Coherent?

Empirical support for any scientific theory comes from observation of things and events in space and time. Hence, if a theory makes no room for such entities (called "local beables"), it might undermine its own empirical basis and thereby face the threat of empirical incoherence. Some authors have argued that this is the case for quantum mechanics, unless we supplement it with local beables at the fundamental level (a so-called "primitive ontology"). The argument involves two premises, namely that (1) quantum mechanics without a primitive ontology has no local beables and that (2) a theory without local beables is empirically incoherent.

The most detailed version of this argument (though without mention of the terms "empirical incoherence" and "primitive ontology") was given by Tim Maudlin (Jnl Phys A 40 (2007), 3151-3172). My paper starts by questioning Maudlin's defence of premise (1). I will then argue that Alyssa Ney's (Synthese online (2015), DOI 10.1007/s11229-014-0633-9) recent response to Maudlin is doubly unjustified, firstly in its sympathy for premise (1), secondly in its rejection of premise (2).

The crucial question behind premise (1) is whether local beables can be derived within a version of quantum mechanics that does not postulate them at the fundamental level. Maudlin (2007, 3161) admits that this might be possible in principle, but he thinks that present attempts to do so lack a clear rationale to regard the derived structure "as physically salient (rather than merely mathematically definable)". In response, Huggett and Wüthrich (Stud Hist Phil Mod Phys 44 (2013), 276-285) point out that physical salience can be assessed "from above", that is, by examining which theoretical structures yield correct empirical predictions. But this does not completely dispel Maudlin's worry. The example he discusses in this context concerns the two different types of local beables that can be associated with the GRW formalism: a matter field (GRWm) or flash-like events (GRWf). Since GRWm and GRWf are empirically equivalent, choosing between them "from above" is impossible. This is a familiar problem of underdetermination, but it is here combined with a less familiar one: not only is the choice between GRWm and GRWf underdetermined by the empirical evidence, but it is also underdetermined by the underlying fundamental theory (GRW without local beables, called GRW0).

However, there is no reason why this twofold underdetermination should be any more worrying than the usual one we face in quantum mechanics anyway. Whoever wants to be a realist about quantum mechanics must opt for one of its versions, based on their non-empirical virtues. This is true for the primitive ontologist (who thinks of GRWm and GRWf in terms of fundamental ontology) as well as for the wave function realist who tries to derive local beables from GRW0. Therefore, insofar as underdetermination does not prevent us from realism about fundamental ontology, it should not prevent us from realism about derivative ontology either.

Still, one might be tempted to endorse premise (1), because the wave function of quantum mechanics does not seem to be the kind of entity from which local beables could emerge. Thus Ney (2015, 15) claims that the wave function could not play the functional role of a three-dimensional object such as a macroscopic pointer. This is directed against the wave function realist's appeal to functionalism, most prominently worked out by David Albert. The curious thing is that Ney (2015, 11) cites Albert's claim that the wave function's dynamics (encoded in the Hamiltonian of the system) "plays the causal role constitutive of there being multiple classical particles in a three-dimensional space", without specifying what is wrong with that claim. But if nothing is wrong with it, then the Hamiltonian may very well be such that these particles form a bound state that constitutes a pointer capable of interacting with other objects in just the way ordinary pointers do. To be sure, it would be hopelessly complicated to actually write down such a Hamiltonian, but this is not a specific problem of wave function realism; it confronts the primitive ontologist in precisely the same way.

Despite her sympathy for premise (1), Ney seeks to defend wave function realism against the charge of empirical incoherence by rejecting premise (2). The claim that empirical coherence presupposes local beables is based on our pre-theoretical beliefs about evidence, and these, she argues, should be replaced by what our best scientific theories tell us about the nature of evidence. But aren't our best scientific theories those which are best supported by empirical evidence? If so, the project of first appraising our scientific theories and then having our beliefs about evidence informed by them is incoherent.

Even setting this problem aside, Ney's (2015, 18) proposed reconceptualization of "evidence" by directly linking the wave function to a state of the world "that is properly described (nonexhaustively) as 'Theorists have acquired evidence for theory T'" does not look promising. It is significantly more problematic than the idea (criticized by Maudlin 2007, 3158-3159) that physical theories should make predictions about our conscious experience. In order to do that, a theory would have to solve the mind-body problem. This would not suffice in the case of Ney's proposal, since there are not even any bodies on her view. What her theory would have to solve is the "mind-wavefunction problem", that is, the challenge of connecting the quantum state of the universe directly to mental states, without passing through the intermediate step of first connecting it to some local beables (pointers, observers' brains etc.), which can then be connected to mental states. Neither Ney nor anyone else has given us any idea how this is supposed to work.

In sum, one can with good reason hold on to premise (2), but should be suspicious of premise (1). The argument from empirical (in-)coherence is therefore inconclusive. In order to save it, the primitive ontologist needs to defend a weakened form of (1), presumably by comparing the explanatory virtues of his approach with the virtues of the attempt to derive local beables from a fundamentally non-local ontology.

Joshua Eisenthal. The Problem of Space

Nowadays it is entirely commonplace to acknowledge a distinction between "pure geometry", as a subfield of mathematics, and "physical geometry", as a subfield of physics. In fact, it would be normal to regard this as just one example of the distinction between pure mathematics and its concrete applications in any of the natural sciences. But this distinction within geometry, between abstract mathematics and descriptions of physical space, is a relatively modern one. After all, geometry (from the Greek meaning ‘to measure the Earth’) originated as an explicit description of spatial relationships, and remained so for most of its history. Indeed, the peculiar nature of classical Euclidean geometry in seeming to combine the rigour and certainty of abstract mathematics with the empirical content of a natural science has impressed and perplexed many great thinkers since it’s conception. As Helmholtz eloquently put the matter:

‘The fact that a science can exist and be developed as has been the case with geometry has always attracted the closest attention among those who are interested in questions relating to the bases of human cognition. Of all branches of human knowledge, there is none which, like it, has sprung as a completely armed Minerva from the head of Jupiter; none before whose death-dealing Aegis doubt and inconsistency have so little dared to raise their eyes. It escapes the tedious and troublesome task of collecting experimental facts... the sole form of its scientific method is deduction. Conclusion is deduced from conclusion, and yet no one of common sense doubts that these geometrical principles must find their practical application in the real world about us.’ (Helmholtz 1870, 'On the Origin and Significance of Geometric Axioms')

Nevertheless, Helmholtz himself – responding to the Kantian account of geometry as a body of synthetic a priori knowledge – was a key figure in the process of refining the distinction between pure and physical geometry. Indeed, Helmholtz was one of the first to recognise that spatial presuppositions had been embedded deep within classical Euclidean geometry, and was one of the trailblazers in untangling such “intuitions” (in Helmholtz’s sense, not Kant’s) from the properly mathematical aspects of the subject.

As geometry has been separated from its moorings in descriptions of physical space, it has become possible to ask new questions. One question is simply: what is the geometrical structure of physical space, given that Euclidean geometry is no longer the answer just by default? A further question is: what are the possible geometrical structures that physical space could have? This latter question is one formulation of the venerable ‘Problem of Space’, a problem tackled in one form of another by Riemann, Helmholtz, Lie, Poincaré, Weyl and Cartan, amongst others. For the purposes of this paper, I will define this as the problem of delimiting the range of candidate physical geometries, i.e. identifying conditions for when a geometrical structure can indeed be thought of as a description of possible spatial structure. I will be concerned in particular with Hermann Weyl’s approach to the Problem of Space immediately following the development of General Relativity. I will aim to show that Weyl’s work constituted an analogous kind of conceptual analysis to Helmholtz’s work fifty years earlier, and will argue that the importance of such work, in both cases, lay in unearthing, or at least suggesting, what it was that contemporaneous physics took space to be like. In Weyl’s case this is of particular relevance because General Relativity remains our best theory of space today.

I will first survey the progress made in the development of the “classical solution” to the Problem of Space, articulated by Helmholtz and other key figures working in the immediate aftermath of the development of non-Euclidean geometries. I will then discuss how this classical solution disintegrated in the upheaval occasioned by the advent of General Relativity, and turn to explore the renewed attempt to tackle the problem by Weyl. However, it is not my aim to defend Weyl’s solution to the Problem of Space here. On the contrary, I believe that a better reason for exploring what Weyl identified as the underlying presuppositions of General Relativity (in particular, the presuppositions regarding physical geometry) is that challenging such presuppositions becomes possible.

My more specific goal will be to demonstrate that a lack of sufficient attention to this question – i.e. to what General Relativity stipulates about the geometrical structure of space – has led to confusion over the status of the metric field. In a relatively recent dispute, some have defended the view that the metric field should no longer be regarded as codifying a property of space itself, but rather regarded as a physical field in space, akin to the electromagnetic field. Others have argued that metrical structure cannot be so separated from a geometrical description of space. Later, I will briefly survey the arguments that have been put forward on both sides. I claim that, in the context of this dispute, it has become clear both that there is no consensus on what the relationship between geometry and physical space can or should be, and that the distinction between pure and physical geometry has become obscure. Thus I hope to demonstrate that one immediate benefit of engaging with the Problem of Space (and Weyl's work in particular) is that we can gain significant insights regarding the status of the metric field in General Relativity.

Samuel Fletcher. Limits of Nagelian Reduction

This presentation concerns the question of whether limiting-type reductions—what Nickles (1973) calls reduction_2—can be accommodated in the Generalized Nagel-Schaffner (GNS) framework (Dizadji-Bahmani et al., 2010) for intertheoretic reduction, which has been argued to avoid most of the problems leveled at the frameworks originally considered by Nagel (1961) and Schaffner (1967). It consists of two parts. The first considers a problem for this accommodation arising from the essentially deductive nature of the relationship that the GNS framework posits between theories: exhibiting a limiting relationship between theories simply does not fit this mold. The second considers a possible response to this problem using a powerful theorem from the theory of uniform spaces. However, I conclude that this offers only a grossly attenuated sense in which limiting-type reductions can be described in the GNS framework.

What's at stake is not merely terminology for describing reduction, or whether a venerable idea for articulating it can be stretched even further to accommodate yet another problematic case. Rather, it is the nature of some reductive intertheoretic relations that is at issue: the GNS framework, even withstanding the modifications that it makes to the frameworks proposed by Nagel and Schaffner, essentially describes a reduction between two theories as a logical relationship: perhaps with the addition of certain bridge rules and similarity relations, one theory's laws are *deduced* from another. Following a natural interpretation of logical deduction, this entails that a GNS reduction involves essentially the containment of one theory, as represented by its laws, within another.

Although some authors (e.g., Butterfield 2011) have asserted that limiting-type reductions fall into this mold, a careful examination of what it takes to make a limiting-type relationship precise reveals significant differences. Limits can be formally defined on a domain by placing a topology on it. One can then interpret the resulting systems of open neighborhoods of the domain's objects as encoding a weak notion of *similarity* amongst them. A sequence of these objects converges to another just in case the elements of the sequence become arbitrarily similar to it. When applied to mathematized scientific theories—the kinds of theories of which it makes sense to take limits—the relevant domain objects are not the laws but the models of the theory. This is one significant difference from the GNS framework. Further, in order to exhibit a (reductive) limiting relationship between theories, the models of both the limiting theory and the limit theory must be described together within the same topological space. Thus there is no sense in which one theory is deduced from another—rather, in limiting-type reduction, both theories are *postulated*. The reduction is also exhibited *relative* to a choice of topology describing a relevant sense of similarity between their models. As an example, I consider the case of the non-relativistic limit of relativity theory.

It is important as well to note that in a limiting-type reduction the similarity relation defined by the topology is not an auxiliary feature, as similarity is in the GNS framework. In limiting-type reductions, the similarity relation does nearly all of the work in drawing a relation between theories, to which the role of deduction of laws is supererogatory. Because it is simply not the case that a limit theory is already contained in its corresponding limiting theory, it would be extremely misleading to continue to insist on describing limiting-type reduction as essentially logical (deductive) in character. This logicality is at the heart of the GNS framework, so a proper understanding of limiting-type reductions seems to challenge it as a truly general framework for understanding intertheoretic reduction.

In the second part of this presentation, I develop and evaluate an interesting possible response to this challenge. There is a kind of structure slightly stronger than topology, called uniform structure, which one can place on a collection of mathematical objects. Just as a topology defines notions of convergence and continuity, uniform structure defines notions of uniform convergence and uniform continuity. Conceptually, while topological structure on a domain defines which mathematical objects in that domain are similar to one another, it does not in general allow for a comparative notion of similarity, e.g., that A is more similar to B than C is to D. Uniform structure adds precisely this comparative notion. There is a remarkable theorem concerning uniform structures, which states that *any* uniform space whatsoever has a unique Hausdorff completion, that is, an extension of the space in which each point is distinguishable, and in which every sequence whose elements eventually get arbitrarily similar to one another converges. Thus, if one equips the models of a limiting scientific theory with a uniform structure, via this theorem one may be able to "deduce" the models of the limit theory.

What chance does the invocation of this theorem have to save limiting-type reductions as Nagelian? There are at least two issues here. The first concerns the nature of the "deduction" provided by the theorem's unique existence claim. One way to explore this issue involves observing that the theorem in question is simply a generalization of the construction of the irrational numbers from Cauchy sequences of rational numbers. Thus, the construction of a limit theory from a limiting theory is a *deduction* from the latter precisely in the sense that one "deduces" the irrational numbers from the rational numbers. Can such a construction, though, truly be called a deduction?

The second issue is that the theorem guarantees a unique completion only relative to the choice of uniform structure on the domain; in general, different choices will lead to different completions. Because there is good reason to believe that at least some theories of interest do not have any canonical notion of similarity on their models (Fletcher 2015), it is doubtful that any choice is determined logically from the theory alone. Thus, even invoking the uniform completion theorem, limiting-type reductions can be said to be Nagelian only in an extremely attenuated sense at best.

Alexander Franklin. Universality Explained?

It is commonly claimed, both by physicists (e.g. Fisher 1998; Kadanoff 2009) and philosophers (e.g. Batterman 2000, 2014; Reutlinger 2014) that the universality of critical phenomena is explained through particular applications of the Renormalisation Group (RG). Such claims are made with a view to using the RG framework to account for multiple realisability more generally, where universality is considered to be a special case thereof.

The details of this explanation are at best spelt out in vague terms, or by reference to paradigm cases. In this paper I argue that the physics underlying such explanations is lacking in important respects.

I take as my model for the explanation on offer recent articles by Robert Batterman in which he gives an abstracted explication of the physical procedures in question. I claim that there are two ways in which this could correspond to the RG derivation of the critical exponents: (i) via a real-space and (ii) via a momentum-space application of the RG.

(i) depends on various extensions of the Ising model which I describe in some detail. These serve as archetypes of the different universality classes. I stress that the derivation does not take diverse systems and justify their inclusion in each universality class, rather universality is assumed and the critical exponents are obtained for each class from its archetype alone. As a case study I consider the inclusion of liquid-gas transitions in the Ising 3D universality class: to the extent that this is physically justified, it does not depend on an RG analysis.

(ii) starts with an effective Hamiltonian which only loosely depends on the details of different physical systems. It can be shown that the addition of various operators to this Hamiltonian would be irrelevant to the derived values of the critical exponents; this implies that multiple Hamiltonians belong to the same universality class. As such, the explanation of universality should involve a justified correspondence between operators and physical systems. I examine the crossover phenomena and demonstrate that operators are only physically interpreted where behaviour is non-universal.

This paper does not aim to denigrate the remarkable achievements of RG physics: for this has been very successful in deriving critical exponents which match those given in experiment. Rather the claim will be that while the RG derivation of the critical exponents is good, the explanation may still be lacking. It is suggested that this is due to the mathematical intractability of physically realistic Hamiltonians.

References

Batterman, Robert W. (2000). 'Multiple Realizability and Universality'. The British Journal for the Philosophy of Science 51.1, pp. 115-145.

Batterman, Robert W. (2014). 'Reduction and Multiple Realizability'. http://philsci-archive.pitt.edu/11190/.

Fisher, Michael E (1998). 'Renormalization Group Theory: Its Basis and Formulation in Statistical Physics?'. Reviews of Modern Physics 70.2, p. 653.

Kadanoff, Leo P et al. (1967). 'Static Phenomena Near Critical Points: Theory and Experiment'. Reviews of Modern Physics 39.2, p. 395.

Reutlinger, Alexander (2014). 'Why Is There Universal Macrobehavior? Renormalization Group Explanation as Noncausal Explanation'. Philosophy of Science 81.5, pp. 1157-1170.

James Fraser. Groundwork for a Neo-Galilean Approach to Idealisation

Idealisations pose a problem for many philosophical accounts of scientific explanation. A prevalent intuition is that explanation is ultimately a matter of pointing to facts about the world which bear some objective explanatory relation to the explanandum. How are we to square this idea with the fact that most, if not all, scientific explanations are based on models that deliberately misrepresent their targets?

One response, most famously associated with McMullin’s (1985) notion of Galilean Idealisation, is that the explanatory capacity of an idealised model can be rationalised by appealing to its relationship to a more realistic ‘de-idealised’ model. In recent years a number of philosophers have argued that, while some idealisations might be accommodated by this approach, the way that idealisations are employed in many scientific explanations is not Galilean in character (Batterman 2009, Morrison 2005, Wayne 2011).

My contention in this paper is that a core claim motivating the notion of Galilean idealisation, which I call the de-idealisation principle, is plausibly true and offers a simple resolution of the apparent tension between orthodox accounts of scientific explanation and the ubiquity of idealisation in scientific practice. The de-idealisation principle says that if an idealised model affords an explanation for why some fact F obtains then F must be derivable from the model and continue to hold when idealised assumptions are replaced by more accurate characterisations of the target. If the simple pendulum model is to explain why a real pendulum has a time period of approximately $2\pi\sqrt{l/g}$, for instance, this fact must be derived from it and continue to hold when realistic air resistive forces are added to the model.

If true, the de-idealisation principle has two immediate payoffs. First, it gives us a way of understanding how idealisations can legitimately feature in explanations. If the explanandum is preserved under de-idealisation then, on any plausible account of explanatory relevance, the features of the target which are distorted by the idealised model must be irrelevant to the question of why the explanandum obtains.

Idealisations do not undermine a models ability to explain then because they misrepresent explanatorily irrelevant features of the target. Second, it allows us to rescue the idea that explanatory power is ultimately derived from worldly facts and dependencies. If the de-idealisation principle holds, it must be features the idealised model shares with a veridical model, and therefore the target system itself, which are responsible for its explanatory success. Once the de-idealisation principle is in place, I suggest, idealised models can be understood as conveying the kind of worldly facts which constitute explanations on a number of specific accounts of scientific explanation.

In the second part of this paper I address some of the criticisms which have been raised against McMullin's notion of Galilean idealisation and argue that they do not give reasons to believe that the de-idealisation approach is violated.

Some of the objections raised in the literature are simply orthogonal to the question of whether the de-idealisation principle holds. Many authors interpret Galilean idealisation as incorporating claims about the epistemic and methodological significance of de-idealisation. The de-idealisation principle does not require that we actually know that the explanandum is preserved under de-idealisation for it to be explanatory however, and I claim that this is not needed to solve the puzzle of how idealised models can explain. Similarly, the de-idealisation principle does not have the implausible methodological implication that idealisations are a temporary measure which will ultimately be eliminated as science progresses. I point out that the idea, pioneered by Strevens (2008), that idealisations can lead to better explanations by facilitating the abstraction of explanatorily relevant information is wholly compatible with the de-idealisation principle.

On the other hand, some putative examples of non-Galilean idealisations do seem to problematise the de-idealisation principle; a prominent case being the thermodynamic limit, in which the volume of a model in statistical mechanics is taken to infinity. I argue, however, that clear cut counterexamples to the de-idealisation principle have not been provided in the literature.

Some authors (Batterman 2009, Wayne 2011) read McMullin as requiring that Galilean idealisations are approximately true, so that infinite idealisations, like the thermodynamic limit, are immediately beyond the scope of his account. I claim that the de-idealisation principle is readily applicable to infinite idealisations; indeed it straightforwardly holds in textbook cases in which the length of a wire is taken to be infinitely long for the purposes of calculating its magnetic field. More problematic for the de-idealisation principle is the singular nature of the thermodynamic limit. When the volume of a model is taken to be infinite discontinuities in macroscopic observables appear which do not occur in any finite model. This will only provide a counterexample to the de-idealisation approach if the novel properties afforded by thermodynamic limit are taken to be physically real features of the target. Batterman (2005), for instance, claims that the discontinuities in the thermodynamic limit correspond to real discontinuities in macroscopic variables that occur under a change of phase. I argue that this is really to beg the question against the de-idealisation principle, in the absence of independent reasons for thinking that phase transitions are physical discontinuities, and suggest that a similar response is available for other putative examples of essential or ineliminable idealisations.

I conclude that the prospects of a neo-Galilean approach to idealisation are better than many authors suppose.

References

Batterman, R. (2005) "Critical phenomena and breaking drops: Infinite idealizations in physics" Studies in History and Philosophy of Modern Physics, 36, 225-244.

Batterman, R. (2009) "Idealization and modelling" Synthese 169.3, 427-446.

McMullin, E. (1985) "Galilean Idealisation" Stud. Hist. Phil. Sci., Vol. 16. No. 3 pp. 247-273

Morrison, M. (2005) “Approximating the Real: The Role of Idealizations in Physical Theory” in Cartwright, N. and Jones, M.R. (ed) Idealization XII: Correcting the Model Idealisation and Abstraction in the Sciences, Rodopi, Amsterdam-New York

Strevens, M. (2008) Depth: an Account of Scientific Explanation. Harvard University Press

Wayne, A (2011) "Expanding the scope of explanatory idealization." Philosophy of Science 78.5, 830-841.

Brian Garvey, ‘The evolution of morality and its rollback’

According to standard Evolutionary Psychology accounts, human moral attitudes are rooted in cognitive modules that are products of evolution in the Stone Age. These modules evolved as adaptive responses to problems of social interaction – e.g.: to reap the benefits of cooperation and exchange, to avoid the costs of being cheated, to know who to help or fight and when. For present purposes, I am willing to grant that this has some plausibility, as an account of what may have happened in the Stone Age. It is likely that early humans had problems of social interaction of the kind described, which created selection pressure. It is unlikely that they could have been solved efficiently by conscious reasoning, but they could have been solved by cognitive modules. If we grant this, it is plausible that cognitive modules of the kind described by Evolutionary Psychology did evolve in the Stone Age.

However, Evolutionary Psychologists further hold that we have those same modules today. “[O]ur moral heuristics are now operating outside the envelope of environments for which they were designed.” [Cosmides and Tooby, 2007] To justify this, they point out that evolutionary change is very slow, and that there is today a general high level of genetic similarity between different human populations all over the world. They also point to similarities between cultures as evidence that there has been no change in underlying mechanisms since the Stone Age. In the present paper I question the claim that such cognitive mechanisms must have remained unchanged since the Stone Age.

To do this, I appeal to the phenomenon of evolutionary rollback (AKA ‘evolutionary streamlining’). This is where an organ becomes non-functional and eventually becomes atrophied or disappears altogether – e.g. cave-dwelling fish losing their eyes, vestigial organs becoming incapable of performing their former functions. This may happen because of the expensiveness of developing an organ that’s not needed. Or there may be costs to having an organ, e.g. risk costs or maintenance costs. Or there may be lack of selection pressure where it is needed to maintain an organ. In this paper, I argue that even if cognitive modules evolved in the Stone Age to solve problems of social interaction, conditions since then have been favourable to the rollback of those modules. This is because of the existence of institutions that solve problems of social interaction, saving us having to. Legal institutions often solve those problems that Evolutionary Psychology says we have cognitive mechanisms to solve. Often, a person who has been mistreated does not even have to make a complaint for a miscreant to be liable for punishment. This can be contrasted with the well-known capuchin monkeys’ responses to unequal treatment: they respond by themselves becoming agitated. Moreover, there is evidence that where external resources are available to perform cognitive tasks, humans often use them in preference over internal ones. (E.g. Sparrow, Liu and Wegner 2011 on the effects of search engines on memory; Tribble 2005 on memory-saving devices in Elizabethan theatre.).

The conjecture proposed here owes much to Rowlands’ (1999) ‘Barking Dog’ principle, and Clark’s ‘007’ principle: “In general, evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand-in for the information-processing operations concerned.” (Clark, 1989.) Shapiro (2010) objects to these principles on the ground that they make Panglossian assumptions. However I argue, first, that anti-Panglossianism would not favour the standard Evolutionary Psychology story. Second, I do not make Panglossian assumptions in the present paper because what I offer, I offer as a conjecture only – with, at the end of the paper, some very brief and tentative suggestions as to how it might be tested.

I will respond to the 'slowness of evolution' argument given by Evolutionary Psychologists in support of their claim that cognitive modules evolved in the Stone Age have not changed. I argue that the kind of evolutionary change that is intrinsically slow is the building-up of complex mechanisms. Undoing that process is much simpler, and can therefore be much quicker. I will also respond to the 'widespread genetic similarity' argument that they have not changed. The rollback of cognitive mechanisms need not require any genetic change: for example in the case described by Sparrow et al. a change from internal to external mechanisms occurred without any genetic change. Moreover, I argue that evidence of cross-cultural similarity does not necessarily support the standard Evolutionary Psychology view, since there is an alternative explanation for such similarity consistent with the conjecture offered here: laws were made in a given society because somebody in that society thought they were a good idea. And sometimes the same thing is a good idea in lots of different cultures.

In the penultimate section, I briefly outline two hypotheses about present-day humans that are consistent with the conjecture. The milder hypothesis is that our moral attitudes and practices are, at least to a greater extent than Evolutionary Psychology supposes, products of our cultural milieu. The more extreme hypothesis is that we don't have, or at least have to a much less extent than is commonly supposed, stable moral attitudes at all: our responses are momentary and are the effects of immediate conditions.

None of this is intended as decisive evidence that rollback of cognitive modules involved in moral attitudes has taken place. I finish by suggesting ways in which we might be able to tell whether attitudes are likely to be rooted in unchanged Stone Age modules or not, given that I have argued that cross-cultural similarity will not do it. I offer two suggestions: (1) We might find good evidence that some attitude or practice is pan-cultural and isn't explainable as obviously a good idea in lots of different contexts. (2) We might find consistent phenomenological and/or behavioural differences between moral attitudes that are indisputably products of culture and ones that are not clearly so.

Alexander Gebharter. Causal exclusion and causal Bayes nets

Causal exclusion arguments, most famously advanced by Kim (2000; 2007), can be used as arguments for epiphenomenalism or as arguments against non-reductive physicalism. Epiphenomenalism is the view that "mental events are caused by physical events in the brain, but have no effects upon any physical events" (Robinson, 2015). Non-reductive physicalism, on the other hand, basically consists of three assumptions: Mental properties supervene on physical properties, mental properties cannot be reduced to physical properties, and mental properties are causally efficacious (cf. Kim, 2007, p. 33).

In a nutshell, exclusion arguments assume non-reductive physicalism and conclude from several premises that mental properties supervening on physical properties cannot cause physical or other mental properties. The notion of causation used in these arguments is, however, typically somewhat vague and not specified in detail. Because of this, the validity of these arguments may depend on the specific theory of causation endorsed (cf. Hitchcock, 2012). In this talk, I reconstruct two variants of the exclusion argument and evaluate their validity within a specific theory of causation, viz. the theory of causal Bayes nets.

The theory of causal Bayes nets (CBNs) evolved from the Bayes net formalism. It was elaborated in detail by researchers such as Pearl (2000) and Spirtes, Glymour, and Scheines (2000). The theory connects causal structures to probability distributions and provides powerful methods for causal discovery, prediction, and testing of causal hypotheses. Furthermore, its core axioms can be justified by an inference to the best explanation of certain statistical phenomena, and several versions of the theory can be proven to have empirical content, by whose means not only the theory's models, but the theory as a whole becomes empirically testable ([anonymized]). So the theory of CBNs probably gives us the best empirical grasp on causation we have so far.

Another strong motivation for this endeavor is that causal exclusion arguments have recently been intensively discussed (cf., e.g., Baumgartner, 2009, 2010; Shapiro, 2010; Shapiro & Sober, 2007; Woodward, 2008) within an interventionist framework of causation a la Woodward (2003), and that interventionist accounts do have a natural counterpart within the theory of CBNs (cf., e.g., [anonymized] or Zhang & Spirtes, 2011). So the hope is that we can draw as of yet unconsidered conclusions for the interventionist debate surrounding causal exclusion arguments from a reconstruction on the basis of the theory of CBNs. This seems especially promising since one of the main problems interventionists have when testing causal efficacy of properties standing in supervenience relationships to other properties is that these properties cannot be simultaneously manipulated by interventions. The theory of CBNs, on the other hand, provides a neat and simple test for causal efficacy not requiring fixability by means of interventions.

The talk is structured as follows: In the first part I briefly introduce two variants of the causal exclusion argument. In the second part, which is also the main part of the talk, I reconstruct these two variants within the theory of CBNs and evaluate their validity. This requires an answer to the question of how supervenience relationships should be represented in CBNs and a test for evaluating whether the instantiation of a property X at least sometimes contributes something to the occurrence of another property Y. I will argue that supervenience relationships can be treated similar to a CBN's causal arrows. A method for testing a property's causal efficacy is already implemented in the productivity condition, which can be proven to be equivalent to one of the theory of CBN's core axioms, viz. the causal minimality condition (cf. Spirtes et al., 2000, p. 31). I conclude by demonstrating that mental properties supervening on physical properties cannot be causally effective if causal relations are assumed to obey the core axioms of the theory of causal nets, i.e., the causal Markov condition and the causal minimality condition. In the third part of the talk I investigate the consequences of these findings for the interventionist debate on the causal exclusion argument.

References

Baumgartner, M. (2009). Interventionist causal exclusion and non-reductive physicalism. International Studies in the Philosophy of Science, 23(2), 161-178. doi:10.1080/02698590903006909

Baumgartner, M. (2010). Interventionism and epiphenomenalism. Canadian Journal of Philosophy, 40(3), 359-383.

Hitchcock, C. (2012). Theories of causation and the causal exclusion argument. Journal of Consciousness Studies, 19(5-6), 40-56.

Kim, J. (2000). Mind in a physical world. MIT Press.

Kim, J. (2007). Physicalism, or something near enough. Princeton University Press. doi:10.2307/j.ctt7snrs

Pearl, J. (2000). Causality (1st ed.). Cambridge: Cambridge University Press.

Robinson, W. (2015). Epiphenomenalism. In E. N. Zalta, The Stanford encyclopedia of philosophy. Retrieved from http://plato.stanford.edu/archives/spr2015/entries/epiphenomenalism/

Shapiro, L. A. (2010). Lessons from causal exclusion. Philosophy and Phenomenological Research, 81(3), 594–604. doi:10.1111/j.1933-1592.2010.00382.x

Shapiro, L. A., & Sober, E. (2007). Epiphenomenalism -- the Do"s and the Don"ts. In G. Wolters & P. Machamer, Studies in causality: Historical and contemporary (pp. 235–264). University of Pittsburgh Press.

Spirtes, P., Glymour, C., & Scheines, R. (2000). Causation, prediction, and search (2nd ed.). MIT Press.

Woodward, J. (2003). Making things happen. Oxford: Oxford University Press.

Woodward, J. (2008). Mental causation and neural mechanisms. In J. Hohwy & J. Kallestrup, Being reduced: New essays on reduction, explanation, and causation (pp. 218–262). Being Reduced. Oxford University Press.

Zhang, J., & Spirtes, P. (2011). Intervention, determinism, and the causal minimality condition. Synthese, 182(3), 335–347. doi:10.1007/s11229-010-9751-1

Marton Gomori, Laszlo E. Szabo and Zalan Gyenis. Operationalist Approach to Quantum Theory: Two Representation Theorems

The paper is a report on the first main results of a larger project aiming to reconstruct the foundations of quantum mechanics in purely operationalist terms. We focus on the presentation of two main representation theorems, without presenting the lengthy and technical proofs.

First we formulate a general framework in operational terms describing a typical experimental scenario in physics: one can perform different measurement operations on a physical system, each of which may have different possible outcomes. Empirical data are the observed relative frequencies of how many times different measurement operations are performed and how many times different outcome events occur, including the joint performances of two or more measurements and the conjunctions of their outcomes.

We do not make a priori assumptions about these relative frequencies. Any truth about them is regarded as empirical fact observed in the experiments. For example, the fact that two measurements cannot be performed simultaneously reveals in the observed fact that their conjunction is always of zero relative frequency. Similarly, the obvious facts that an outcome event cannot occur without the performance of the corresponding measurement operation, that two different outcomes of the same measurement cannot occur simultaneously, but one of them necessarily occurs whenever the measurement operation is performed, etc., also reveal in the observed relative frequencies. In terms of the observed frequencies, we will specify a few less obvious features of a general experimental setup, concerning the possible influence of one measurement operation on the outcomes of another measurement.

In the first representation theorem we prove that whatever the physical system in question is---traditionally categorized as classical or quantum---everything that can be meaningfully described in empirical/operational terms can be described within the classical Kolmogorovian probability theory. In other words, we give a proof of the Kolmogorovian Censorship Hypothesis (Szabó 1995, Bana Durt 1997, Szabó 2001, Rédei 2010) within our general operational framework.

The relative frequencies of measurement operations constitute the “soft” part of the empirical data, as they may depend on circumstances outside of the physical system under consideration; for example, on the autonomous choice of a human. One can hope a scientific description of the system only if the two things can be separated. We will present the conditions when this separation is possible. Given that these conditions are satisfied empirically, we define a concept within the Kolmogorovian model that can be interpreted as characterization of the system's state, in the sense that it characterizes the system's future probabilistic behavior against all possible measurement operations of the experimenter.

From mathematical point of view, the state so defined is a correlation vector (Pitowsky 1989) constructed from conditional probabilities in the Kolmogorovian model. Due to the fact that the conditional probabilities in question generally belong to different conditions, the state of the system does not belong to the classical correlation polytope---no matter if the physical system in question is traditionally categorized as classical or quantum. In some situations however the physical system and the measurement operations are such that the resulted conditional probabilities constitute a classical correlation vector. One can show that in this case the system's state admits property/elements of reality interpretation.

In a second representation theorem, we show that everything that can be meaningfully described in empirical/operational terms can also be represented in the Hilbert space quantum mechanical formalism. There always exists:

(1) a suitable Hilbert space

(2) such that the outcomes of each measurement can be represented by orthogonal projectors,

(3) the states of the system can be represented by suitable density operators,

(4) and the probabilities of the measurement outcomes can be reproduced by the usual trace formula of quantum mechanics.

Moreover, in the case of real-valued quantities (if the measurement outcomes are “coordinatized” by real numbers),

(5) each quantity can be associated with a suitable self-adjoint operator,

(6) and the expectation value can be reproduced by the usual trace formula applied to the self-adjoint operator.

Thus, in some sense, the second representation theorem can be interpreted that the basic premises of the quantum mechanical formalism are the expression of the fact that the system can be described in empirical/operational terms.

References

Bana, G. and Durt, T. (1997): Proof of Kolmogorovian Censorship, Foundations of Physics 27, 1355.

Pitowksy, I. (1989): Quantum Probability -- Quantum Logic, Lecture Notes in Physics, vol. 321., Springer, Berlin.

Rédei M. (2010): Kolmogorovian Censorship Hypothesis for general quantum probability theories, Manuscrito -- Revista Internacional de Filosofia 33, 365.

Szabó, L. E. (1995): Is quantum mechanics compatible with a deterministic universe? Two interpretations of quantum probabilities, Foundations of Physics Letters 8, 421.

Szabó, L. E. (2001): Critical reflections on quantum probability theory, in: John von Neumann and the Foundations of Quantum Physics, M. Rédei and M. Stoeltzner (eds.), Kluwer Academic Publishers, Dordrecht.

Rosa Hardt. The Interdependence of Emotion and Sensory Experience

On the face of it, emotional experience and sensory experience seem distinct. Prinz’s (2004) theory of emotions as embodied appraisals preserves this pretheoretical intuition. On his account, emotions are simultaneously bodily experiences and appraisals of a creature’s relationship to the world. They do not include sensory experience as a constituent.

I am in agreement with Prinz that emotions are felt states of the body that express how we find ourselves in the world. However, I claim that separating emotional and sensory experience neither accurately reflects our phenomenology nor finds support in recent work in neuroscience and clinical psychology. Taking evidence from these fields into account should lead us to view emotional experience and sensory experience as intimately intertwined. My aim is to demonstrate why such a position not only better fits the available evidence than Prinz's theory, but also meets further criteria for a good scientific theory: it is consistent, parsimonious, and fruitful.

On phenomenological grounds, Merleau-Ponty (1996) argues that we perceive the world before us when there is coherence between (i) the information we receive from the world and (ii) our bodily engagement with the world. Our capacity for bodily engagement is tied to our capacity for experiencing the world emotionally – as bearing on us in a certain way. And experiencing the world as bearing on us in a certain way incorporates the sensory experience of the object towards which our emotion is directed. Thus sensory experience depends upon emotion. Phenomenology therefore lends support to the claim that sensory and emotional experience are interdependent.

Barrett and Bar (2009) draw a similar conclusion from neuroscientific evidence. The lateral orbitofrontal cortex (OFC) is involved in the conscious perception of objects and this process involves integrating information from the body and the world to form a judgment on the value of a given situation. The OFC is involved in an emotional appraisal of what an object means for us only in virtue of its receiving sensory information from the world. This suggests that sensory information is necessary for emotional experience and therefore leads again to the conclusion that sensory and emotional experience are interdependent.

Evidence from clinical psychology further corroborates this finding. Firstly, people with affective disorders (in particular depression) are found to have altered sensory experiences of the world: it often appears gloomier, more distant, and detached (Ratcliffe, 2013). Secondly, it has been observed that psychosis – a disorder of perception and thought – co-occurs with disruptions in emotional experience. Both findings strongly suggest that alterations in emotion are associated with alterations in perception, and vice versa. Thus again we find a close interdependence between sensory and emotional experience.

The theory that emotions and sensory experience are interdependent is superior to Prinz's in that it accurately predicts the phenomenological, neuroscientific and psychological evidence. Further explanatory virtues also accrue. The theory is consistent: it is not confined to only one domain of inquiry but rather bridges at least three, all of which have different methodologies. It is parsimonious insofar as it posits overlapping mechanisms for the existence of sensory and emotional experience. And it allows us makes novel predictions. For example, it predicts that disturbances to sensory experience will co-emerge with disturbances to emotion.

References

Prinz, J. J. (2004). Gut reactions: A perceptual theory of emotion. Oxford University Press.

Merleau-Ponty, M. (1996). Phenomenology of perception. Motilal Banarsidass Publishe.

Barrett, L. F., & Bar, M. (2009). See it with feeling: affective predictions during object perception. Philosophical Transactions of the Royal Society B: Biological Sciences, 364(1521), 1325-1334.

Ratcliffe, M. (2013). Depression and the phenomenology of free will. Oxford Handbook of Philosophy and Psychiatry, 574-591.

Casey Helgeson. Low confidence in extreme probabilities

Assessment reports from the Intergovernmental Panel on Climate Change (IPCC) comprehensively summarize and evaluate the published scientific literature about climate change and the prospects for mitigation and adaptation. These reports communicate to policymakers the state of scientific knowledge, with a strong emphasis on conveying, for each reported finding, a degree of certainty judged appropriate in light of the available evidence. Since the third assessment cycle (the most recent set of reports concluded the fifth) IPCC authors have approached the task of characterizing uncertainty with the benefit of official guidance notes on the treatment of uncertainties in IPCC reports (Mastrandrea et al., 2010).

These guidance notes prescribe the use of two uncertainty metrics: probability and confidence. Confidence is qualitative, and meant to express the level of scientific understanding that backs up a given finding; it is assessed at one of five levels, from 'very low' to 'very high'. The two metrics are often used together, for example: 'In the Northern Hemisphere, 1983–2012 was likely the warmest 30-year period of the last 1400 years (medium confidence)'. Here a probabilistic statement ('likely' has an official translation of '66–100% chance') is further qualified by a level of confidence (in this case, medium).

How best to use this two-metric framework to characterize and communicate uncertainties (and how to improve the framework itself) is an ongoing topic of debate (Adler and Hadorn, 2014; Yohe and Oppenheimer, 2011). One puzzle concerns how to interpret extreme probabilities (close to 1 or to 0) qualified by low confidence. Kandlikar et al. (2005) and Risbey and Kandlikar (2007) argue that such combinations are incoherent: 'It wouldn't make any sense to declare that an event was extremely likely and then turn around and say that we had low confidence in that statement. For example, if we declare that it is extremely likely to rain tomorrow, but then say that we have very low confidence in that statement, that would lead to a state of confusion'.

I argue that low confidence in extreme probabilities is NOT incoherent, and that these authors' examples exploit a subtle and illicit shift between confidence in (1) the actual occurrence of the event in question and (2) attribution of a probability to that event occurring.

Consider another example (a toy example — but the point is conceptual, not scientific). An urn contains 100 balls, each either red or blue. Suppose you caught a glimpse of the balls being poured into the urn from a sack. You saw lots of red and perhaps a bit of blue. Asked about the proportion of red balls you respond that you really couldn't say. When pressed, your best guess is 98 red, 2 blue. This is equivalent to saying that the chance of drawing red is 98%. Your confidence in this statement, however, is low. I suggest there is nothing incoherent in this instance of low confidence in an extreme probability. 98 red (low confidence) is no more or less problematic than 72 red, or 44 red, or any other number (all with low confidence).

But there is another possible source of incoherence, or at least confusion: not low confidence in HIGH probabilities, but low confidence in very PRECISE probabilities. Unlike the Risbey and Kandlikar (2007) examples — which use precise probabilities — most IPCC findings use probability ranges. The menu of calibrated probability language on which IPCC authors draw includes a range of terms, and the terms that indicate more extreme probabilities indicate, at the same time, more precise probabilities. For example: 'likely' means probability .66-1.0; 'very likely' means .9-1.0; 'virtually certain' means .99-1.0. As the probabilities get higher, they also get sharper.

Consumers of scientific information implicitly assume that greater precision indicates greater confidence. This makes a claim like 'such and such event has 99-100% chance to occur (low confidence)' seem confusing. But the problematic aspect is due to the precision of the probability, not its extremeness (closeness to one). To modify the Risbey and Kandlikar (2007) example from above: If we declare that the chance of rain tomorrow is [exactly 57%], but then say that we have very low confidence in that statement, that would lead to a state of confusion.

But while reporting low confidence in very precise probabilities has the potential to confuse, it is not, strictly speaking, incoherent. Indeed, the most developed framework for systematically relating probability to confidence, and managing trade-offs between the two (Hill, 2013), entails that low-confidence probabilities are always sharper than their high-confidence counterparts. This raises a conflict between on the one hand, what is formally the most coherent and systematic approach to probability and confidence, and on the other hand, what is most intelligible by typical readers of IPCC assessment reports.

References

Adler, C. E. and G. H. Hadorn (2014). The IPCC and treatment of uncertainties: Topics and sources of dissensus. Wiley Interdisciplinary Reviews: Climate Change.

Hill, B. (2013). Confidence and decision. Games and Economic Behavior 82, 675–692.

Kandlikar, M., J. Risbey, and S. Dessai (2005). Representing and communicating deep uncertainty in climate-change assessments. Comptes Rendus Geoscience 337(4), 443– 455.

Mastrandrea, M. D., C. B. Field, T. F. Stocker, O. Edenhofer, K. L. Ebi, D. J. Frame, H. Held, E. Kriegler, K. J. Mach, P. R. Matschoss, G.-K. Plattner, G. W. Yohe, and F. W. . . Zwiers (2010). Guidance note for lead authors of the IPCC Fifth Assessment Report on consistent treatment of uncertainties. Intergovernmental Panel on Climate Change (IPCC).

Risbey, J. S. and M. Kandlikar (2007). Expressions of likelihood and confidence in the IPCC uncertainty assessment process. Climatic Change 85: 19–31.

Yohe, G and M. Oppenheimer (2011). Evaluation, characterization, and communication of uncertainty by the intergovernmental panel on climate change—an introductory essay. Climatic Change 108: 629–639

Elselijn Kingma. Functions at the interface of biology and technology: synthetic biology, health and disease

Synthetic biology is the designing and building of new biological parts and processes. This in principle allows for the production of completely human-intended, purposefully designed biological organisms. It can be tempting to think of such an organism as organic or biological machines. This places synthetic biology at the interface of the biological and the engineering world – and makes it an interesting perspective from which to reconsider existing philosophical analyses of function.

It is uncontroversial that functional analyses are appropriate in both engineered systems and organisms, but they are analysed differently in each domain; biological functions are often explicitly analysed in terms of natural selected effects, whereas engineering functions often appear the designer's intent or human use. In this context, synthetic biological organisms appear to present a problem: they are not the product of natural selection, so how can they possess biological functions?

In this paper I analyse function judgments in synthetic biological organisms and compare them to cultivated and co-evolved organisms. I argue, first, that functional analysis in artifacts and organisms is far more continuous than one might presume; we can and should bridge the gap that has opened up between biological and technological function. Second, I shall argue that the aetiological analysis of biological function need to be interpreted more flexible than is usually proposed; in a way that encompasses selective and reproductive processes other than natural selection. Moreover I shall argue that agriculture and domesticated animals provide us with reasons for doing this independent from synthetic biology.

Karen Kovaka. Rejecting Replicators

The argument of this paper is that it is wrong to think of replicators as the basic unit of biological inheritance. When someone asks what kind of hereditary material is transmitted through biological inheritance processes, we should not answer in terms of replicators. I begin with a characterization of the extended replicator view. Then I register an objection to one particular element of the view: the claim that in order to count as hereditary material, an object must be subject to cumulative selection. I provide a conceptual reason to doubt this claim, and then offer empirical counterexamples to it, focusing on phenotypic plasticity and genetic accommodation in threespine stickleback fish. I discuss evidence for the idea that non-replicator hereditary material is important in evolution and argue that continued empirical work on this idea will benefit from a rejection of the replicator view.

Debate about the replicator/interactor distinction has raged among evolutionary biologists for several decades, and while my interest in replicators is very much informed by the existing literature, I reorient the familiar debate in two ways. First, rather than conflating genes and replicators, I focus on the idea of the extended replicator because it accommodates both genetic and non-genetic replicators in its view of inheritance. Second, I meet the replicator view on its home turf by focusing on explaining evolvability, rather than development. My goal is not merely to show that non-replicator hereditary material exists and is relevant for explaining parent—offspring resemblances, but also that non-replicators are a source of heritable variation and may be crucial to explaining the evolvability of lineages.

The replicator view of inheritance is the dominant view of inheritance among scientists and philosophers. On this view, a replicator is an object that both copies itself and makes a causal contribution to an organism's phenotype because it has been selected to make that causal contribution. Genes are the paradigmatic replicators, but burrows, nests, and symbiotic bacteria also qualify. The argument of the replicator view is that even though non-replicator developmental resources such as nutrition are causally important to phenotypic development, they are not subject to cumulative selection like replicators are. Evolutionary complexity requires cumulative selection, so only replicators are important for explaining the evolvability of complex lineages.

My challenge to the replicator view begins by examining the view's efforts to take non-genetic inheritance seriously. Proponents of the replicator view want to capture something more general than genes; they want to capture all of the hereditary material that actually matters for explaining the evolution of complex living systems. Genes may form a supermajority in the world of biological inheritance, but a full explanation of evolved complexity will not limit itself to the most common kind of hereditary material. Given that this is the goal of the replicator view, an important question for the view is whether its general characterization of the units of inheritance does actually capture all of hereditary material that needs to be included in an explanation of evolved complexity.

I argue that we should be suspicious of the replicator view's claim to capture all of the relevant hereditary material.

First, I identify a point of disagreement with the replicator view. I accept the claims that replication is required for cumulative selection and that cumulative selection is required for complex systems to evolve. But it does not follow from these two claims alone that replicators are the only type of hereditary material can be important in the evolution of complex systems. Proponents of the replicator view need to defend a further claim, that an object must be subject to cumulative selection in order to count as hereditary material that matters for evolvability. This is what I do not accept.

Second, I explain why I do not accept this element of the replicator view. Cumulative selection is only one part of the evolvability story. Another part of the story is the generation of variation. Often, variation is generated by genetic mutations, which are immediately subject to selection. But we know that in other cases, variation is generated by plastic phenotypic responses to novel environmental conditions. If this variation can later become subject to selection, then it, too, can matter in evolution, even when replicators are not responsible for the initial appearance of the phenotypic variation.

Third, I consider the empirical evidence supporting the idea that non-replicators can generate variation and that such variation can later become subject to cumulative selection. Many biologists have dedicated themselves to the investigating the possibility that novel traits can emerge because of phenotypic plasticity and non-replicator environmental conditions, and I survey the state of the science on this question. I also discuss recent empirical work on genetic assimilation, a process by which variation due to plasticity can become subject to cumulative selection. I focus on plasticity in threespine stickleback fish, which are emerging as a model system for these issues.

The empirical evidence justifies us in believing that non-replicator hereditary material plausibly has been and still is important in evolution. Nonetheless, much of the evidence for the importance of non-replicator genetic material is piecemeal. There are many demonstrations of individual parts of the hypothesized process, but arguably no confirmed instances of macroevolutionary change in a natural population due novel environmental conditions, phenotypic plasticity, and genetic assimilation.

This evidential situation means that the grounds for rejecting the replicator view are not purely empirical. Rather, I argue that the replicator view itself discourages the research that could provide additional insight into non-replicator hereditary material and its role in evolution. It is this dampening effect of the replicator view that justifies us in rejecting of the view. In order to promote a better investigation into the nature and role of hereditary material in evolution, we should adopt a broader account of biological inheritance, rather than limiting the scope of hereditary material to replicators.

References

Bateson, P. and Gluckman, P. (2011). Plasticity, robustness, development and evolution. Cambridge, UK: Cambridge University Press.

Calcott, B. and Sterelny, K. (Eds.) (2011). The major transitions in evolution revisited. Cambridge, MA: MIT Press.

Mameli M (2004) Nongenetic selection and nongenetic inheritance. Brit J Phil Sci 55, 35–71.

Pigliucci, M. (2001). Phenotypic plasticity: beyond nature and nurture. Baltimore, MD: Johns Hopkins University Press.

Schlichting, C.D. and Wund, M.A. (2014). Phenotypic plasticity and epigenetic marking: an assessment of evidence for genetic accommodation. Evolution 68-3, 656-672.

Sterelny, K. (2011). Darwinian spaces: Peter Godfrey-Smith on selection and evolution. Biology and Philosophy 26, 489–500.

Sterelny, K., Smith, K.C and Dickison, M.: 1996, ‘The Extended Replicator’, Biology and Philosophy 11, 377–403.

Raymond Lal. The topology of contextuality: a unifying concept in quantum theory and logic

1. Background: non-locality and contextuality in quantum theory.

Quantum theory is our most successful theory of physics. It also bears many counter-intuitive features. A list of them would include the existence of entangled states, the no-cloning theorem, quantum teleportation, and many others. These features suggest a quite different view of reality from classical physics. This is most clearly shown by the famous 'no-go theorems' of quantum mechanics. The most intensely studied of these are Bell's theorem [1] and the Kochen-Specker theorem [2], which have been respectively used to demonstrate the non-locality and contextuality of quantum theory.

The body of work around these results has led to two long-standing views in the foundations and philosophy of physics which we shall challenge: one methodological and one philosophical. First, in methodological terms, the study of quantum contextuality has largely been carried out in a concrete, example-driven fashion, which makes it appear highly specific to quantum mechanics. Second, the philosophical lesson of quantum no-go theorems has tended to be in terms of properties that quantum theory lacks---e.g.causally explicable correlations; measurement results that are independent of the experimental setup---especially in comparison to classical physics. More generally, no-go results constrain any attempt to replace quantum mechanics with a theory that has a similar structure to classical physics. Since the latter is the gold-standard for a realist theory of physics, many have taken no-go theorems to place quantum mechanics in severe opposition to realism.

Taken together, these challenges lead to my main claim: we should reconfigure our attitude towards the role of contextuality in quantum theory: rather than an attack on realism, it is best understood as a pervasive feature of fundamental results in foundational fields.

2. A methodological challenge: unifying non-locality and contextuality.

The first step in this endeavour is to recognise that non-locality and contextuality are examples of the same type of phenomenon. The first hints of this appeared in the work by Fine [3], who showed that non-locality in the bipartite Bell setup corresponds to what we might call 'generalised contextuality'. More specifically, the usual definition of non-locality involves providing a common cause explanation of four 'partial' probability distributions P(a_x,b_y), one for each pair of measurements. Fine showed that non-locality is equivalent to fact that there is no global probability distribution on all four measurements, P(a_1,a_2,b_1,b_2), that marginalises to the empirical probabilities P(a_x,b_y).

Roughly speaking, we can say that non-locality demonstrates a conflict between partial data and global data.

In fact, Fine's theorem can be generalised to all non-locality proofs; but also to standard proofs of contextuality. Hence the notion of generalised contextuality subsumes the property of non-locality and standard contextuality.

The definition of generalised contextuality makes no reference to the mathematics of quantum theory, e.g. operators on a Hilbert spaces---it is 'theory-independent'. This leaves open the possibility that it could be applied in other domains. In fact, recent work by the present author and collaborators has shown that contextuality is a general and indeed pervasive phenomenon, which can be found in many areas of classical computation, such as databases and constraints. So the notion of generalised contextuality not only unifies different quantum phenomena, but it also unifies quantum phenomena with concepts from different scientific fields.

3. A foundational challenge: contextuality as topological twisting.

A particularly interesting example of this concerns logical paradoxes.

The classic Liar sentence, ``This sentence is false'', has been generalised to a family of paradoxes known as 'Liar cycles'. A Liar cycle of length N is a sequence of statements:

S_1: S_2 is true,

S_2: S_3 is true, ...

S_{N-1}: S_N is true,

S_N: S_1 is false.

Any subset of up to n-1 of these equations is consistent; while the whole set is inconsistent---once again, a conflict between partial data and global data. Indeed, this is more than an informal analogy: the present author and collaborators have shown that Liar cycles are examples of generalised contextuality. Further, we have shown that, when generalised contextuality is formalised as a kind of discrete fibre bundle, it corresponds to 'topological twisting'---analogous to the non-orientability of the Moebius strip.

4. A feature not a bug?

This work offers the possibility of quantifying contextuality as the presence of an extra property as compared to classical physics, rather than as the 'lack' of a certain classical model. In consonance, Howard et al. [4] have shown that contextuality is responsible for the possible advantages of quantum computation over classical computation.

So, whilst philosophers have traditionally viewed contextuality as an obstruction to realism, physicists and computer scientists have discovered that quantum contextuality is a resource. Moreover, the presence of contextuality in other fields suggests that its philosophical significance may lie in general aspects of reasoning, rather than the more narrow ontological concerns of physics. A general mathematical theory is required to fully understand this phenomenon, and I will argue that the topological tools that we have identified provide the clues to constructing this theory.

References

[1] John S. Bell. "On the Einstein-Podolsky-Rosen paradox." Physics, 1(3):195–200, 1964.

[2] Simon Kochen and Ernst P. Specker. "The problem of hidden variables in quantum mechanics."" Journal of Mathematics and Mechanics, 17(1):59–87, 1967.

[3] Arthur Fine. "Hidden variables, joint probability, and the Bell inequalities."" Physical Review Letters, 48(5):291, 1982.

[4] Mark Howard, Joel Wallman, Victor Veitch, and Joseph Emerson. "Contextuality supplies the 'magic' for quantum computation."" Nature, 510(7505):351–355, 2014.

Jürgen Landes and Jon Williamson. How an objective Bayesian integrates data

The Wider Picture

Computers have made it possible to collect and store large data sets. Reasoners would like to make use of as many data sets as possible which are as large as possible while still allowing for computationally feasible inferences. Ideally, one could simply combine all the available datasets.

These days, several datasets involving hundreds of variables and thousands of observations are routinely collected in many applications. Unfortunately, different datasets tend to measure different variables, even when the datasets are collected with same application in mind. For instance, it is common in systems medicine to have datasets measuring proteomics, transcriptomics, metabolomics, clinical data, and patient-reported outcomes, and for these datasets to have very few variables in common. How do we integrate all this data?

One approach to data integration is motivated by objective Bayesian epistemology (OBE), which holds that a rational agent ought to adopt as a representation of her degrees of belief the probability function with maximum entropy, P, from all those calibrated to her evidence [3].

In this talk I shall assume that the agent’s body of evidence consists of a collection of datasets and nothing else. Furthermore, I assume that the datasets are large and reliable enough that each dataset distribution provides an accurate estimate of the frequency distribution of the measured variables, and that they are consistent in the sense that these marginal frequency distributions are satisfiable by some joint probability function defined on the set V of all the variables measured by the datasets. The agent’s credence function P will be defined on this larger set V of variables. OBE holds that P should be calibrated to each marginal distribution of observed frequencies, i.e., P should agree with each dataset distribution.

In general, finding the function which has maximum entropy is a computationally hard optimisation problem. In this talk I show how, in a wide range of cases, one can compute P without optimising at all, via a Bayesian net representation of P. A Bayesian net representation of the credence function P which is motivated by OBE is called an objective Bayesian net [2].

To the extent that the dataset distributions only approximate the corresponding marginal frequency distributions, the function P that we determine here should be thought of as an approximation of the credence function warranted by OBE.

Bayesian Networks

Bayesian networks can often be used to efficiently represent and reason with joint probability distributions. The machine learning community has developed efficient algorithms to learn a Bayesian network from a dataset; see, e.g., [1]. We shall take such an algorithm as given, and apply it to each dataset—i.e., use it to learn a Bayesian net Bi which represents the marginal frequency distribution P*i determined by dataset DSi over its set Vi1 of variables. Each such Bayesian net consists of a directed acyclic graph on the set Vi of vertices together with the dataset distribution of each variable conditional on its parents in the directed acyclic graph.

Two Datasets

The first result I will present is a simple way to obtain P in the case of two datasets. One first learns Bayesian networks B1 and B2. After some manipulation of these Bayesian nets one analytically obtains P without solving any optimisation problem at all.

Centred Datasets

The above recipe for two datasets can be generalised to a collection of datasets DS1,...,DSh, in case the collection of datasets is centre. A collection of datasets DS1,...,DSh is called centred, if there exists a dataset DSm such that every variable which is measured in more than one dataset is also measured in DSm.

I will show how to analytically obtain P for a collection of centred datasets without solving any optimisation problem at all. Time permitting, I shall touch on how to obtain P in more complicated cases.

References

[1] Ioannis Tsamardinos, Laura E. Brown, and Constantin F. Aliferis. The max-min hill-climbing Bayesian network structure learning algorithm. Machine Learning, 65(1):31–78, 2006.

[2] Jon Williamson. Bayesian Nets and Causality. Oxford University Press, 2005.

[3] Jon Williamson. In Defence of Objective Bayesianism. Oxford University Press, 2010.

Dennis Lehmkuhl. The neighborhood of General Relativity in the space of (spacetime?) theories

How ‘special’ is General Relativity (GR) as compared to other theories? The answer to this question depends on what other theories we compare GR to: other field theories or just other spacetime theories? I will argue that Einstein himself saw GR not primarily as a theory of spacetime, but as a field theory unifying gravity and inertia. I will then show that his interpretation of GR as a unification of gravity and inertia is only possible because of the way the different fields couple in GR, and compare GR to a much later theory (Jordan’s theory from the 1950s, the first scalar-tensor theory). The comparison will show that it is the coupling structure that ensures the motion of particles on geodesics, and thus the possibility for Einstein to interpret the theory as a unified field theory (of gravity and inertia).

Owen Maroney and Daniel Bedingham. A flash, a collapse and a boundary condition: where did the asymmetry come from?

The world is full of processes that are frequently observed happening in one direction, but never observed in the reverse: an ice cube in a glass of warm water will melt and merge with the water, but ice cubes do not spontaneously form, freezing out of a glass of warm water. The asymmetry in time of these phenomena seems puzzling in light of the fact that our most fundamental theories of physics do not show any such asymmetry. According to the equations of motion, if a process in one direction is allowed, then the reverse process should also be allowed. The traditional explanation for this is in terms of allowing solutions to the equations of motion to be constrained with special initial conditions, but leaving final conditions unconstrained. Ultimately this solution is applied to the whole universe, which is then assumed to start in a special state characterised by extremely low thermodynamic entropy.

For those that are dissatisfied with this, a way out might be to modify the equations of motion, introducing a fundamental time asymmetry. Collapse models of quantum theory, independently introduced as a means for solving the quantum measurement problem, appear to provide just such a modification. In a collapse model, the wavefunction is treated as representing the physically real properties of systems. To avoid the conceptual problems associated with representing macroscopic objects by quantum superpositions, the dynamics of quantum systems are modified, so that superpositions collapse well before reaching macroscopic proportions. This modification appears to introduce a time reversal asymmetry into the dynamics of the wavefunction, since the collapses affect only the future state of the system, not the past.

We will challenge whether collapse models can really account for macroscopic time asymmetry in this way. There are grounds to be suspicious. The physical collapse of wavefunctions is a model for the Born rule of quantum theory. However, the Born rule can be stated in a time symmetric manner, using the Aharanov-Bergmann-Lebowitz rule (or its generalisation to the consistent histories or the two state vector formalisms). The usual way in which quantum mechanics is applied involves the construction of ensembles of pre-selected states. This allows us to make predictions about the future state. Time-symmetric formulations of quantum theory show that by post-selecting the final state we can also make retrodictions using the same laws, implying that the basic laws of quantum mechanics can be understood in a time symmetric way. Any time asymmetry is a result of constraining our statistical ensembles with initial conditions, but leaving final conditions unconstrained.

We analyse three cases where the addition of a physical collapse process to quantum theory appears to have introduced a fundamental time asymmetry into the dynamics. Firstly, collapses affect the future evolution of the system, but not the past. During the collapse process an initially dispersed wavefunction might become spontaneously localised about some position. The time reverse of this process, an initially localised state instantaneously dispersing, is not an allowed solution. Secondly, when viewed as a process of dispersion, the dynamical equations do not show time reversal symmetry. There is a correlation between changes in position and changes in momentum that is not invariant under a time reversal symmetry. Finally, wavefunction collapse models show a monotonic increase in mean energy over time.

Our analysis will make use of the `flash ontology' for collapse models. Collapses are localised events, in space and time, that occur randomly with a probability given by the Born rule. A flash ontology treats the locations of the collapses themselves as the basis for the local beables of the theory - the mathematical counterparts to real world events. On a fine-grained scale the world appears as if composed from many discrete points. The local density of these points give a representation of the location of matter. The role of the wavefunction is purely nomological, determining the probabilities for the various collapse locations. This means the wavefunction can be relegated to the initial time, from which it does not need to evolve, with the collapsing evolution corresponding to an updating (conditioned on the history of realised collapses) of the rule for determining the probability of future collapses. This implies that even the initial wavefunction can be replaced by a sufficiently long period of collapse data. The evolving wavefunction is then just a convenient calculation tool for making the theory Markovian, and only the flashes remain.

On this basis of these ideas, we show that, given a valid set of collapse centres, we can form a picture of a collapsing wavefunction evolving either forward or backward in time, and in each case, the locations of subsequent collapses will satisfy the Born rule. In this sense a collapse model can have time reversal symmetry. The forward going and backward going pictures will not be the same for the same set of collapse data. In particular the backward going wavefunction will be affected by collapse events in the future and not the past. Nonetheless, the collapse centres, the local beables of the theory, will be consistent between the two pictures. For examples of the first two cases we show that the fixed collapses which were generated by the forward in time dynamics are distributed as though they had been generated by the backward in time collapse dynamics. In the final case we will show how the apparent monotonic increase in energy is compatible with time symmetric dynamical laws, and arises through the asymmetric use of initial and final conditions. Any physically observed time asymmetries that arise in such models are due to the asymmetric imposition of initial or final time boundary conditions, rather than from an inherent asymmetry in the dynamical law. This is the standard explanation of time asymmetric behaviour resulting from time symmetric laws.

Kerry McKenzie. Intrinsicality and The Goldilocks Principle: Fundamentality as an Untapped Resource for Structuralism

Right from its inception, structuralists have identified structural features with ‘those that are not intrinsic’ (Maxwell 1970, 188; see also Ladyman and Ross 2007, 151). Since ontic structural realism (OSR) is the thesis that ‘relational structure is ontologically fundamental’ (ibid. p145), OSRSists must deny that there are any intrinsic properties at the fundamental level. As such, mass, charge, or whatever properties will turn out to define the fundamental kinds in nature must be shown to be extrinsic if OSR is to go through. Unfortunately, however, metaphysicians of science tend to be united in their intuition – despite being divided on almost every other issue – that this is not the case (see e.g. Chakravartty 2012, 204; Bird 2007, 125). The question then is how the structuralist should proceed in the face of this scepticism from the rest of the field.

Not surprisingly, structuralists have not been willing to submit to force majeure with respect to their opponent’s intuitions. Quite the contrary, as committed naturalists they have objected that whatever people's intuitions are on matters of fundamental ontology are entirely irrelevant to them: it is only the relevant theories of physics that can be invoked to settle such matters (Ladyman and Ross op cit, passim). However, this move faces the problem that we do not presently possess a physics theory that we regard as truly fundamental: as such, we lack precisely the theory suitable for consulting to adjudicate the issue. It seems, therefore, that the debate over the presence or absence of intrinsic properties in the fundamental base is simply at an impasse.

In this paper, I argue that there is a route out of this quandary that has not been exploited until now. This route proceeds from fundamentality considerations in relativistic quantum physics, and settles the matter in OSR's favour. This is because these considerations show that all fundamental properties, *whatever* they may turn out to be, are extrinsic in character.

At the heart of my proposal are two fairly uncontroversial claims. The first is that, while it is true that we currently lack a fundamental theory, we do have a framework for thinking about theories that may well be fundamental. This is the framework of quantum field theory (QFT). The second is that the only way to define a fundamental kind in the framework of QFT is as a kind featuring in a fundamental theory. Since owing to the possibility of the creation and decay in this context, mereological conceptions of fundamental kindhood have long been thought to be ruled out (see e.g. Heisenberg 1975; Weinberg 1996). Since we do, however, have precise criteria of what a fundamental theory is in the context of QFT, this offers an alternative route for defining fundamental kinds.

It is these criteria just mentioned applying to fundamental theories in QFT that lie front and centre in my argument. The crucial fact is that, while it is true that we do not yet know the identity of the fundamental theory, we do know that this framework places tight constraints on any theory that deserves to be called 'fundamental'. Since a fundamental theory is one that never needs to be replaced no matter how high the energy of interactions grow, any theory of QFT that deserves to be called fundamental must stay consistent in the limit of infinite energy (owing to the continuity of Minkowski spacetime). At present, the only theories satisfying this stringent requirement – at least that are mathematically tractable – are the `asymptotically free theories'. It has been known since the work of Coleman, Gross and Wilczek in the early 1970s that these theories must be renormalizable local gauge theories (see Wilczek 1997 for discussion). This insight shows that fundamental properties cannot be regarded as intrinsic for two reasons.

The first reason is that the requirement of gauge symmetry implies that fermions must be accompanied by gauge bosons if they are to feature in fundamental theories. Since intrinsic properties are those that may be possessed by an entity independently of what the rest of the world is like (see e.g. Dunn 1990, Weatherson and Marshall 2012), this implies that the defining properties of fundamental fermions cannot be regarded as intrinsic properties. (To make it explicit: fundamental fermions are those that feature in fundamental theories; without gauge bosons, there cannot be fundamental theories; therefore whatever properties define fundamental fermionic kinds cannot be possessed independently of what the world at like, and as such are not intrinsic.)

The second reason is that one can show that the property of asymptotic freedom is preserved only if tight constraints on the number of fermions featuring in the theory are respected. As a QFT textbook puts it, ‘theories of non-Abelian gauge fields and fermion multiplets are asymptotically free only if the theory does not have too many fermions’ (1998, p. 278). By the same reasoning as above, then, we must therefore say that the very existence of a fundamental kinds is contingent on there not being too many other kinds co-existing with the first. As such, the properties of definitive of that kind again cannot be regarded as intrinsic. While a less obvious route to extrinsicality than the last, this approach is even more extensive in its reach, for it also has implications for the properties of bosons. Throw too many fermions in this mix, and we are left with no fundamental bosons at all: as such, the properties defining fundamental bosons cannot be regarded as intrinsic any more than those of fermions can.

In sum, then, I will that fundamentality constraints in QFT suggest that a sort of ‘Goldilock’s principle’ governs the fundamental kinds in nature: a kind can only qualify as fundamental if there are other kinds inhabiting the world alongside it – just so long as there are not too many. It follows from this that the fundamental kind properties simply cannot be regarded as intrinsic. Rather than being made up of entities with their own intrinsic natures, then, the inhabitants of the fundamental level are entities that must be highly sensitive to the totality of what else exists in order to be inhabitants of that level at all. And since OSR is the thesis that fundamental entities of the world lack intrinsic properties, it follows that this thesis is left standing tall.

Michael Miller. Exact Models and Physical Semantics

In her recent book, Interpreting Quantum Theories, Laura Ruetsche explains that "Given a theory T , . . . we confront the exemplary interpretive question of how exactly to establish a correspondence between T 's models and worlds possible according to T "(Ruetsche 2011, pp. 102-103). It is standardly assumed that to address this interpretive question one associates physical semantic content with the elements of the mathematical structure of the theory. When T is taken to be quantum field theory, the interpreter's project is complicated by the fact that the referent of T is ambiguous between the perturbative formalism used by many physicists and a collection of axiomatic approaches pursued predominantly by mathematical physicists. Interpreting quantum field theory requires clarifying the nature of the relationship between these formalisms.

The perturbative expansion for empirically adequate models of quantum field theory is likely only asymptotic to exact models of the axioms that are used to specify the non-perturbative content of the theory. They are asymptotic expansions that do not satisfy the strong asymptotic condition and thus do not contain sufficient information to reconstruct an exact model of the non-perturbative content of the theory.1 This does not make such expansions unrigorous as is commonly assumed in the philosophical literature. Rather, I argue that they hold information about the world in a different way than is typically assumed in accounts of mathematical representation. Moreover, I argue that the interpretation of quantum field theory should not proceed exclusively from an axiomatic articulation of the theory and the exact models of those axioms alone for precisely this reason. To do so is to press the axiomatic approach into a service for which is was not designed2, and to make it impossible to use the information about the world provided by the asymptotic expansions of empirically adequate models. The positive part of the paper shows how a physical semantics can be assigned in the case of models that are not captured exactly by available non-pertubative characterizations of the theory.

There are a number of extant accounts of the role of physical axiomatization which provide important background to my argument.3 However, I will show that the nature of the relationship between empirically adequate models and axiomatic formulations of the theory necessitates an account of the role of axiomatization that is distinct from the one commonly assumed in the philosophy of quantum field theory, and the philosophy of physics more generally. Developing such a positive account requires explicating the role that the axioms play in establishing a physical semantics for the theory.4 I argue that the empirical success of an asymptotic expansion does not warrant inference to the representational success of any particular mathematical structure as a given expansion can be asymptotic to a potentially infinite class of structures. I argue that the axioms can function as a guide for how to attribute a physical semantics to all of those structures.

In order to motivate the need for such an enhanced account of the attribution of physical semantics, I illustrate the account with two additional examples. In particular, I consider simple systems in both classical mechanics and non-relativistic quantum mechanics. These cases have the advantage that what count as the basic postulates of the theories are relatively well agreed upon. In each case, one obtains divergent asymptotic expansions for important physical observables. By applying the proposal for the assignment of physical content outlined above to show that it successfully accounts for these cases as well. In this sense, the argument of the paper is intended to be relevant to the interpretation of physical theories in general, and not just quantum field theory in particular.

Notes

1 For further details see, for example, (Reed and Simon 1978).

2 In other work I am providing a historical analysis of the motivations underlying the development of the general theory of quantized fields. In particular I demonstrate that one of the central goals of the program was to extract from quantum electrodynamics a set of physical postulates that any theory of quantized fields should satisfy, and express them in mathematically well defined terms.

3 See, for example, (Corry 1997; Fraser 2011; Redei 2014; Sauer 2002; Stoltzner 2001; Stoltzner 2001; Wallace 2011).

4 This is the subject of both (Wilson 2014) and (Curiel 2014). While I agree with these authors that the proposals along the lines of (Beth 1960) and (van Fraassen 1970) that have dominated the literature for decades are deficient, their positive proposals are quite different from the one I advocate.

References

Beth, E. W. (1960). Semantics of physical theories. Synthese 12 (2-3), 172-175.

Corry, L. (1997). David Hilbert and the axiomatization of physics (1894-1905). Archive for History of Exact Sciences 51 (2), 83-198.

Curiel, E. (2014). On the Propriety of Physical Theories as a Basis for Their Semantics. Unpublished Manuscript .

Fraser, D. (2011, May). How to take particle physics seriously: A further defence of axiomatic quantum field theory. Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics 42 (2), 126-135.

R'edei, M. (2014, May). Hilbert's 6th Problem and Axiomatic Quantum Field Theory. Perspectives on Science 22 (1), 80-97.

Reed, M. and B. Simon (1978). Methods of Modern Mathematical Physics Volume 4: Analysis of Operators. Academic Press.

Ruetsche, L. (2011). Interpreting Quantum Theories. Oxford University Press.

Sauer, T. (2002). Hopes and Disappointments in Hilbert's Axiomatic "Foundations of Physics". In M. Heidelberger and F. Stadler (Eds.), History of Philosophy of Science: New Trends and Perspectives, pp. 225-237. Kluwer Academic Publishers.

Stöltzner, M. (2001). Opportunistic Axiomatics - von Neumann on the Methodology of Mathematical Physics. In M. Redei and M. Stoltzner (Eds.), John von Neuman and the Foundations of Quantum Physics, pp. 35-62. Kluwer Academic Publishers.

van Fraassen, B. C. (1970). On the extension of Beth's semantics of physical theories. Philosophy of Science 37 (3), 325-339.

Wallace, D. (2011, May). Taking particle physics seriously: A critique of the algebraic approach to quantum field theory. Studies in History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics 42 (2), 116-125.

Wilson, M. (2014). A Second Pilgrim's Progress. Unpublished Manuscript .

Thomas Moller-Nielsen. Weak Discernibility, Again

According to one recent influential construal of the Principle of the Identity of Indiscernibles (PII) proposed by Simon Saunders (2003, 2006) (and previously defended by Quine 1976), elementary bosons violate the PII, but fermions (and composite bosons with fermionic constituents) do not: for the latter, but not the former, are always at least *weakly discernible* --- they invariably stand in at least some symmetric but irreflexive physical relation. The conclusion drawn by Saunders is that fermions' weak discernibility in turn guarantees their status as "objects" in some appropriate sense, whereas elementary bosons' failure to stand in such symmetric but irreflexive relations reveals that they are not to be construed as "objects", but rather merely as "mode[s] of the corresponding quantum field" (2006, p. 60).

As an illustrative example, consider the spherically-symmetric singlet state of two intrinsically identical fermions. Despite this state's very high degree of symmetry, the fermions in question will nevertheless still stand in the irreflexive relation "... has opposite direction of each component of spin to ...". It is the fact that fermions *always* stand in at least *some* such irreflexive relation to others that is said to ensure their status as "objects". The same cannot be said, however, of the elementary bosons: it is possible for them all to exist in exactly the same quantum state such that none of them stands in even an irreflexive physical relation to any other. Their status as "objects" is therefore, according to Saunders, not always guaranteed.

We might summarise the general argument as follows:

(P1) Weak discernibility is a necessary and sufficient condition for objecthood.

(P2) Fermions are always at least weakly discernible; the elementary bosons are not.

(C) Fermions are invariably objects; the elementary bosons are not.

Let us grant the second premise: fermions, but not bosons, are always at least weakly discernible. (But see Muller & Seevinck (2009), who argue that bosons can be weakly discerned as well.) But what about the first premise? Why think that weak discernibility is a necessary and sufficient condition for objecthood?

The precise content of this question, unfortunately, is somewhat obscured by the fact that there is no real consensus among philosophers as to what the appropriate criteria are for "objecthood" are: indeed, one might be tempted to read (P1) as having the status of a mere *definition* of the word "object" (i.e. as being, in some minimal sense, an "entity" which is always at least weakly discernible). Most theorists, however, take (P1) to have an implied content which is much more substantive: more specifically, they view Saunders as claiming that fermions' weak discernibility *grounds*, or "metaphysically explains", their status as numerically distinct entities, and, hence, as "objects". (In this context theorists usually implicitly assume that numerical distinctness is a necessary and sufficient condition for objecthood.) Moreover, many of these theorists have criticised this proposal for its implicitly involving a dubious circularity: weak discernibility, they claim, rather than grounding the numerical distinctness of the relevant relata, illicitly presupposes it.

Steven French (2014, p. 40) has recently put this worry as follows:

"... [A] circularity threatens: in order to appeal to such [irreflexive] relations, one has had to already individuate the particles which are so related and the numerical diversity of the particles has been presupposed by the relation which hence cannot account for it."

As French notes, this worry basically reprises a dispute which is very old --- centuries, if not more --- in the history of philosophy: namely, that of whether relations *in general* are capable of accounting for the numerical diversity of their relata. Many philosophers who take a position on this issue today are liable to claim that those who disagree are in thrall to mere metaphysical prejudice; at best, it is a dispute unlikely to be resolved any time soon.

What many of these contemporary philosophers seem to have missed --- or, perhaps, ignored --- is that the notion of weak discernibility was never originally intended to have any specific impact on this debate. As I shall argue, Saunders' main motivation in (re-)introducing the term to philosophy was *methodological*, rather than metaphysical. In other words, weak discernibility's intended role was to serve as an essential aspect of a broader "logical aid" for *interpreting* physical theories; it was not intended to have any particular bearing on the more robustly metaphysical question of whether relations, irreflexive or otherwise, are capable of grounding numerical diversity.

That Saunders sees weak discernibility as primarily serving a methodological as opposed to metaphysical function, and that he does not construe objects' weak discernibility as being that which serves to ground their numerical diversity, is, I take it, reasonably clear from what he writes in his original (2003) paper:

"I do not suppose there is anything wrong with identity, taken in an irreducible sense; whatever objects there are, we know what the identity relation is among them; given objects, identity can look after itself. [...] The proposal, rather, is that in a situation in which we *do not know* what physical objects there are, but only, in the first instance, predicates and terms, and connections between them, then we should tailor our ontology to fit; we should admit no more as entities than are required that can be made out by their means." (Saunders 2003, p. 292.)

As I understand Saunders' proposal, then, weak discernibility's primary function would appear to be to serve as part of a general method of "extracting" or "reading off" objects --- and, in particular, *talk* of objects, using declarative sentences and standard first-order predicate logic (2003, p. 290) --- from physical theory. The obtaining of the relevant irreflexive physical relations is therefore meant to serve as the minimum (and sufficient) condition for when one may permissibly speak of there *being* objects in the appropriate sense; moreover, one should refrain from granting objecthood to those (putative) entities whose (alleged) numerical distinctness cannot be specified using the predicates and terms drawn from physical theory alone.

The interpretative recipe that Saunders seems to be suggesting might therefore be usefully summarised as follows:

1. Begin with an "initially interpreted" physical theory T. This will include various interpreted physical predicates, terms, and *putative* objects.

2. See which putative objects are at least *weakly discernible* according to the various physical predicates that appear in T (stripped of identity).

3. Take the putative objects that can be suitably discerned in this way to constitute (what we might call) T's *genuine ontology*; putative objects which cannot be so discerned are thus not part of T's genuine ontology.

It should be quite clear, then, that the methodological and metaphysical construals of weak discernibility are orthogonal to one another: for accepting --- or rejecting --- this *interpretative* package simply does not bear on the separate *metaphysical* question of what, precisely, grounds facts about numerical diversity. The interpretational program that Saunders is plausibly advocating is therefore logically distinct from the debate which appears to have consumed the majority of philosophers when it comes to discussions of weak discernibility: moreover, it is a program which, in my view, finally deserves our scrutiny.

Thus, in summary, this paper essentially serves a dual purpose: first, as an attempt to rehabilitate discussion of this "methodological" construal of weak discernibility; and second, as an attempt to make some tentative inroads in assessing its overall tenability.

Robert Northcott. Approximate truth and scientific realism

Historically, the motivation for defining a scientific theory's approximate truth has mainly come from the scientific realism debate. Indeed, finding such a definition has been seen by some as essential for buttressing the realist position. As anti-realists often point out, philosophers have had great difficulty in giving a plausible and consistent account of approximate truth. Yet a good and useful definition of it can be found nevertheless – but only once we cast off this inherited entanglement with scientific realism. It turns out that influential recent work in the causation literature is a much more fertile inspiration, as approximate truth can be well defined in causal terms. The crucial move is to change our focus from theories as a whole instead to application-specific models.

Why reject the attempt to define approximate truth for theories as a whole? One reason is the crude induction that not many have been convinced by the results so far. But more analytically, consider a fundamental difficulty facing any such attempt: namely that a theory's errors can be very empirically costly in one application but not at all costly in another, thus leaving it ill-defined how serious those errors are in any context-independent or absolute sense. For example, in dynamical systems theories, should we prefer a theory whose dynamic equations are almost correct but whose empirical predictions quickly become wildly wrong, or a rival theory with the opposite pattern? The answer is inevitably interest-relative or application-specific. Moreover, there seems to be no good way of making sense of 'ontological' approximate truth independent of empirical success in particular applications.

Another common difficulty is the attempt to capture in one measure both accuracy and comprehensiveness, so that, for instance, we are not forced to rank a trivial tautology above the false but widely useful Newtonian theory. But again this difficulty melts away once we relativise approximate truth to specific applications, for then Newtonian theory's very wide scope will immediately be reflected by it scoring well in many different applications.

Even if we accept these general reasons for preferring an application-specific approach, how exactly should such an approach be carried out? Drawing on previous work I show how a definition can be framed in terms of causes. Roughly speaking, according to it a model is approximately true if it captures accurately the strengths of the causes actually present in a given target situation. Accordingly, getting closer to the truth consists in capturing these causal strengths more accurately. In order to make this idea precise, the notion of causal strength, or degree of causal importance, has to be defined, and then also a measure of closeness between a model's assignment of causal strengths and the true assignment. Completing these tasks leads naturally to a definition with just the application-specificity needed to solve the problems above that confound a more generalist approach.

The definition brings with it other advantages too. One is that it is not an abstract logical measure but rather is couched in the causal language that actual scientists use. Moreover, a high score for approximate truth now guarantees empirical success. A high score also carries another easily interpreted implication, namely that it guarantees accurate quantitative predictions of the impact of interventions – here, the recent extensive literature connecting causation with interventions pays dividends. Moreover, the counterfactual element of causation allows us to avoid rewarding 'fluke' empirical successes. (The definition also has several other attractive technical features besides.)

There is another way too in which the application-specific approach ties neatly into scientific practice. Much recent philosophical work has focused on the many cases in which progress does not come via development of new theory but rather via a lot of case-specific extra-theoretical investigation. Often, we end up with an empirical model tailored very closely to a unique event or task, but which cannot be derived from theory or even piecemeal from a group of theories or by trial-and-error tinkering with a theory's parameter values. Progress towards the truth in such cases is well represented by an application-specific causal definition – but is inevitably invisible to definitions of approximate truth in terms of theories as a whole.

This, finally, is where the inheritance from the scientific realism debate reveals itself to be unhelpful. In particular, that debate has usually concerned itself with whether we should be realist about theories. 'Convergent realism', for instance, postulates that our best theories are over time gradually getting closer and closer to the truth. Yet one implication of viewing approximate truth application-specifically is that progress towards the truth is only ever a local not a global phenomenon – quite contrary to convergent realism. Historically, from Popper on, the approximate truth literature has overwhelmingly been focused on theories not application-specific models, reflecting its roots in the realism debate. Yet, I argue, this has unfortunately led us away from a definition that can actually work.

Flavia Padovani. Coordination, Measurement, and the Problem of Representation of Physical Quantities

A condition for the objectivity of scientific knowledge rests on the ability to coherently represent the behaviour of measured objects as a good approximation of a theoretical ideal, which appears as some form of "natural prior" with respect to actual measurements. Measurement outcomes can be inferred from instrument indications only against the background of an idealised model, which strictly depends on the scientific theory in use. What one obtains is thus a construct, rather than a "brute fact". Furthermore, the parameters that appear in scientific theories and equations are not pre-existing quantities. As the history of science illustrates, simple items of our scientific knowledge that we take for granted actually arise as outstanding achievement of our scientific conceptualisation and technical progress. In fact, the individuation of certain quantities as parameters for the relevant laws and equations is often developed together with the instruments required in order to measure them.

In his Scientific Representation (2008), van Fraassen has emphasised how measuring should be considered as a form of representation. In fact, every measurement identifies its target in accordance with specific operational rules within an already-constituted theoretical space, in which conceptual interconnections can be represented. So, this space provides the range of possible features related to the measured items expressed within the language of the relevant theory. Without this space of pre-ordered possibilities no objects of representations can be given. In this sense, the act of measuring is "constitutive" of the measured quantities as it allows for the coordination of mathematical quantities to elements of reality, thereby providing meaning to the abstract representations through which we seek to capture physical phenomena.

In recent years, there has been a revived interest in the notion of "coordination" especially in relation to the issue of scientific representation as van Fraassen has described it in his (2008, ch. 5). In this connection, Hans Reichenbach's original account of coordination has revealed to be particularly interesting. In his early work, however, the idea of "coordination" was employed not only to indicate a class of very general, theory-specific fundamental principles to be potentially revised (or relativized) in the passage to a new scientific theory—as is usually emphasised in the literature on coordination—but also to refer to a number of other "more basic" principles. In Reichenbach's early work, these "basic" principles are related not much to the structural features of a theory, but rather to the conceptual presuppositions required in order to approach the world through measurement in the first instance. Those basic principles are primarily necessary to translate the unshaped material from perception into some quantities that can be used within the mathematical language of physics. Quite interestingly, in his early writings many of these coordinating principles are conceived as preconditions both of the individuation of physical magnitudes and of their measurement. In other words, they are not limited to the definition of quantity terms but they also involve the individuation of what these quantity terms are supposed to be coordinated to.

The aim of this paper is to reassess Reichenbach's approach to coordination in light of recent literature on measurement and scientific representation.

Veli-Pekka Parkkinen. Mechanism-based extrapolation reconsidered

Biomedical research often relies on evidence obtained in animal models to establish causal claims about humans. This practice faces the problem of extrapolation: how to justify a causal claim about some target system based on evidence obtained in an experimental system that is causally dissimilar to the target.

In this presentation I consider a recent account by Daniel Steel that claims that extrapolation of causal claims can be justified by articulating the mechanisms that support the causal relation of interest. I argue that this account fails to provide robust justification for extrapolation. It relies on an assumption that mechanisms exhibit modularity: it must be possible to change features of individual components without changing the causal structure of the rest of the mechanism. However, it is unlikely that this assumption can be justified by any general argument, for two related reasons. Firstly, the modularity-assumptions required for particular extrapolation tasks will vary depending on the specification of the causal claim of interest. For example, extrapolating quantitative aspects of a causal relation requires more demanding modularity-assumptions than extrapolating a qualitative causal claim. Secondly, when studying a causal relation in an experimental context, variation in the mediating mechanism's components and its causal context is deliberately controlled for. As a consequence, the specific modularity-properties of the relevant mechanisms are often unknown.

I conclude that mechanistic reasoning can ameliorate the problem of extrapolation, but provides no general warrant for it. However, once we acknowledge the assumptions that go into mechanistic reasoning, the warrant that mechanistic knowledge gives to extrapolation can be evaluated case-by-case.

Viorel Paslaru. Integrative Pluralism in Scientific Explanations, and a Lesson from Ecology

Integration of diverse explanatorily relevant information plays a pivotal role in the recent accounts of explanation in biology due to Carl Craver (2007, 2006; Kaplan and Craver 2011) and Ingo Brigandt (2010, 2013). While they agree that mechanistic explanation requires multifield and multilevel integration, they disagree on the role of mathematical models. Using examples from neuroscience, Craver argues that mathematical models describe and predict, but do not explain. Brigandt scrutinizes cases from evolutionary developmental and systems biology and argues that mathematical models can be indispensable in biological explanations. Since their findings are intended to apply to biology in general and possibly beyond it, and to examine whether and how they need to be amended, I test them by confronting their ideas with research from ecology on a representative case: competition. Like many accounts of ecological phenomena, explanation of competition is multifield, multilevel and uses mathematical models. In light of this examination, and given an argument based on the current understanding of ecology, I argue for the integrative nature of ecological explanation.

Extant research on completion in ecology typically explores one or two of its aspects, but not all of them at once: (a) population dynamics and their models, (b) mechanisms of competition, and (c) causal relations among competing individuals, or populations and other factors influencing competition. Lotka-Volterra model of competition is a classic example of using models to examine population dynamics. The work of David Tilman on models of competition between populations for one and several limited resources and his exploration of mechanisms that drive the population level dynamics exemplifies research on both populations and mechanisms. Recent work by Eric Lamb on how causal interactions between competing populations of plants and causal models of those interactions can be investigated using structural equation modeling exemplifies the focus on causal relations. The disparate lines of research raise the issue of integrating them into a coherent conception. I address this issue and show that description of mechanisms plays a central role, for they underlie population dynamics, as well as causal relations captured by structural equation modeling.

Mathematical models used to model population dynamics are explanatory in two ways: (i) they offer explanations of population dynamics that is clearer than the detailed tracking of specific interactions between individuals, and (ii) in some explanations it is the mathematical model, rather than the biological component that does all the explanatory work (Colyvan, 2008). Given this, Craver's view has to be amended to account for the indispensable role of mathematical models in ecological explanations. His view has to admit that some mathematical models are explanatory of ecological phenomena in virtue of the mathematical component and because they summarize features that characterize numerous individuals, while others offer satisfactory explanations of population dynamics, but which would be further increased if it were coupled with accounts of mechanisms underlying the population dynamics. It is this idea that guided the research program of Tilman to link mechanisms to population dynamics. Moreover, part of population ecology is also the program of identifying laws of ecology, such as pursued by Weber (1999), Colyvan and Ginzburg (2003b, a) Lange (2005, 2000). Yet while their arguments convincingly support the thesis that there are laws of ecology, they are insufficient to show that nomothetic explanations are satisfactory, for they do not take into account the mechanistic details of central interest to ecologists. In fact, Weber, Colvyan and Ginzburg admit that laws are just one component of ecological explanations, while descriptions of mechanisms are the other one. Exclusive focus on mechanisms, as argued for by Machamer, Craver, Darden (Machamer et al., 2000; Craver, 2005, 2006; Darden, 2013), Bechtel (2005), and others, is not sufficient either, for it lacks the resources for quantitative prediction and population modeling, and so it has to be complemented by mathematical models, which vindicates the views of both Brigandt and Carver. Use of structural equation modeling to elicit causal relations and produce causal models is another research program in which mathematical models play a necessary explanatory role, contrary to Craver. Although causal models offer satisfactory explanations, ecologists consider necessary for the purpose of offering a more complete explanation to also examine the mechanisms that underlie the causal relations and the changes in population dynamics.

In light of the foregoing, I show that the view of Brigandt on the nature of explanation is more adequate to account for the explanatory practice of ecologists that relies heavily on mathematical models in building explanations. To account for this aspect, Craver's conception would have to renounce its ontic commitment and admit that some explanations and mathematical models are explanatory although they do not describe causal relations.

I end with an additional argument in favor of integrative pluralism as the adequate characteristic of explanation in ecology and in support of commitment to integrative pluralism of Brigandt and Craver. I argue that problems pertaining to the nature of explanation in ecology have to be addressed against the backdrop of current understanding of what is ecology: "the scientific study of the interactions that determine the distribution and abundance of organisms" (Krebs 2009, 5). I show that Krebs' characterization of ecology implies a normative account of explanation in ecology that requires integration of epistemic units or explanations that deal with population dynamics, characteristics individuals as representative individuals of species involved, as well as interactions between individuals of different species. I also show that taken separately, none of the three approaches to study competition offers an explanation that meets the implicit normative view, but offer explanations of partial value. Instead, it is their integration that meets the implicit norms on explanation.

Tomasz Placek. Indeterminism and bifurcating geodesics

Last decades have seen many constructions of non-isometric extensions of maximal globally hyperbolic spacetimes of general relativity (GR)— see Ringstr\"om (2009) for a review. Such extensions are rather particular: They are not globally hyperbolic (which is implied, at least in the vacuum case, by the Choquet-Bruhat and Geroch theorem). Closed timeline curves are ubiquitous In new regions of the extensions. The union of such extensions is a manifold, but it is non-Hausdorff. And (typically) non-isomeric extensions are rare, i.e., they arise only for a "small" subset of an allowable range of some parameter. Accordingly, if being rare is synonymous to non-generic, the existence of non-isometric extensions accords nicely with (a version of) the Strong Cosmic Censorship Conjecture, according to which for generic initial data to Einstein's equations, the maximal globally hyperbolic development has no extension. However, unless being rare is equated with unphysical (which is unwarranted), non-isometric maximal extensions M1 and M2 of a vacuum GR spacetime M provide evidence for indeterminism of GR, in the sense of Butterfield's (1989 p. 9) influential definition of determinism: the image of M in M1 is isomorphic to the image of M in M2, whereas M1 and M2 are not isomorphic (as there are no fields in these spacetimes, isometry coincides with isomorphism).

Yet, non-isometric extensions of a GR spacetime have a feature that cast shadow on the above verdict of indeterminism: although an initial spacetime is extended in nonequivalent ways, no geodesics in the initial spacetime bifurcates into (the new regions) of separate extensions. In other words, the non-Hausdorff manifold resulting from pasting together of non-isometric extensions satisfies the existence and uniqueness property for geodesics: given a metric of an appropriate continuity, a point, and a vector at this point, there is a unique maximal geodesics that passes through this point and whose tangent at this point coincides with the vector. This feature has been welcome by physicists who feared that bifurcating paths means an observer's world-line undergoing splitting in one spacetime (see e.g., Ellis and Hawking, 1973, p. 174). It was also used to adjudicate between acceptable non-Hausdorff spacetimes, i.e., those without bifurcating geodesics, and unacceptable ones, i.e., involving bifurcating geodesics (Geroch 1968).

Non-isometric extensions of a GR spacetime are naturally seen as alternative possible developments of this spacetime. Accordingly, a non-Hausdorff manifold that the union of extensions gives rise to can be perceived as a modal structure encompassing alternative possible spacetimes rather then a single spacetime. Maximal Hausdorff sub-manifolds of this manifold are then interpreted as (GR) spacetimes (for more on this view, see M\"uller 2013). Note that on this modal reading, bifurcating world-lines are to be expected in a non-Hausdorff manifold, as they indicate an object with more than one possible path, each possible path continuing through a different (Hausdorff) spacetime. Thus, from the modal perspective the uniqueness property of geodesics looks weird: it suggests a picture of a universe with alternative possible developments, but with every object in this universe having a unique possible evolution.

It is instructive to see how non-Hausdorffness of a manifold resulting from gluing together non-isometric extensions combines with non-bifurcating geodesics. A maximal development of maximal globally hyperbolic spacetime is typically constructed by taking the quotient of some base manifold with respect an equivalence relation (the technique is used for instance to produce extensions of Misner's spacetime). This procedure has three remarkable features: First, a set of increasingly closer paths in the initial manifold is identified by the equivalence relation as a single geodesic in the resulting manifold. Second, some two arbitrarily close points in the initial spacetimes induce equivalence classes (points in the resulting manifold) that cannot be separated by non-overlapping neighborhoods in the resulting manifold (these points are witnesses of non-Hausdorffness of this manifold). Finally, in the resulting manifold each point belongs to exactly one geodesics. The result is the following behavior of geodesics in the non-Hausdorff manifold: some two geodesics follows their separate paths in the region that is the image of the initial spacetime (one that got multiply extended). On the border of this region and the "new" regions (i.e, added by the extensions), the two geodesics come arbitrarily close to each other, as they pass through topologically inseparable points (i.e., points that witness a failure of Hausdorff property). But immediately affect this close encounter they go apart. My diagnosis of this peculiar behavior is that it has not anything to do with the determinism question; in particular, Hausdorffness of a manifold should not be used for a criterion of whether this manifold represents a single spacetime or many possible spacetimes (contra Muller, op. cit.).

The contrast between non-isometric extensions and non-bifurcating geodesics merits a distinction between the received notion of determinism that focuses on an entire universe (a spacetime, or a system conceived of as a separate possible world), and an individuals-based notion of determinism. This latter notion begins with asking what individuals (local, persisting in time objects) are admitted by a theory, possibly under a given interpretation. The theory (under a given interpretation) is then classified as deterministic if every individual it admits has, according to this theory, a unique possible evolution. After developing this notion semi-formally, I will illustrate how it analyzes cases discussed in the debate over determinism of physical theories.

References

Butterfield, J. (1989).The hole truth. Brit. Jour. Phil. Sci. 40(1):1--28.

Geroch R., (1968). Local characterization of singularities in general relativity. Jour. Math. Phys. 9, 450

Hawking, S.W. and Ellis, G.F.R. (1973) The large scale structure of space-time. Cambridge UP.

Müller, T. (2013). A generalized manifold topology for branching space-times. Philosophy of Science, 80:1089–1100.

Ringström, H. (2009). The Cauchy Problem in General Relativity. EMS Publishing House, Zürich.

Bryan W. Roberts. Geometrizing Quantum Theory

Felix Bloch pointed out what is now a well-known way to view the states of certain quantum systems geometrically. Suppose we are dealing with the absolute simplest quantum system, a motionless particle with spin. States on the same ray are taken to be equivalent. Basis vectors |0> and |1> are identified with a pair of poles on the sphere. Then an arbitrary pure state psi can identified with a point, psi = cos(x/2)|0> + exp(iy)sin(x/2)|1>. The result is the famous Bloch sphere.

The question that we are concerned with is, how far can this geometrization of quantum theory be pushed? The pure states of the Bloch system can be represented quite literally as the surface of a manifold. But more is known. The mixed states turn out to be represented by the interior points of the sphere. And the transition probabilities turn out to be given by the relation |<psi1,psi2>| = cos d, where d is the shortest path along the great circle connecting the points on the sphere corresponding to these states. The question thus arises: how many of the facts about a quantum system can be expressed using intrinsically geometric structures?

The answer we will argue for is, *all of them*. In particular, every quantum system admits a geometric representation in exactly the same sense that it represents a Hilbert space representation, and all of the same measurable facts are captured.

This work develops a proposal originally due to Kibble (1979), who showed that one can view the set of rays of any Hilbert space as a manifold, and the inner product as inducing a symplectic form and a metric. It was later shown that a host of further quantum structures can be expressed in this formalism, such as the observables, transition probabilities, and uncertainty relations; these were discussed by Gibbons (1992), Brody et al. (1998) and Ashtekar and Schilling (1999). These efforts have not just been in the interest of mathematical exercise: the geometric framework has been put to serious physical and philosophical use. Ashtekar and Schilling found used it as part of a search for new physics, in allowing a precise framework for systematically relaxing the assumptions of ordinary quantum theory. The framework has also been used by philosophers to study the generality of no-cloning results (Teh 2012), as well as the nature of quantum clocks and time observables (Roberts 2014).

However, character of many of the technical results relating ordinary and geometric quantum theory is to identify a number of features of familiar quantum theory within the context of a geometric theory. In this paper, we will make this equivalence considerably more general and more robust: a quantum theory (relativistic or non-relativistic) always admits a geometric representation.

We begin by pointing out that quantum theory, viewed most generally as a C* algebra, is deeply related to another algebraic structure, which \citet{landsman1998mathematical} has called a Jordan-Lie-Banach (JLB) algebra. We show that the relationship between them is equivalence in a very strong sense: C* algebras are related to JLB algebras by an isomorphism of categories. We then state and prove a representation theorem, that for every JLB algebra there exists a geometric quantum system (in particular, a Kahler manifold) that expresses all its algebraic and probabilistic facts. Our construction is very much analogous to the GNS-theorem, which constructs a Hilbert space representation from a C* algebra in a similar way.

The result is a precise illustration of the sense in which quantum theory is equivalent to a geometric theory: it always admits a geometric representation as a Kahler manifold.

Davide Romano. The Meaning of the Mass in Bohm's Theory and Classical Mechanics. A Case Study from the Classical Limit.

The experiments with neutron interferometry performed by Colella et alii (1975) and Staudenmann et alii (1980), mainly discussed in the philosophical literature by Brown et alii (1995) (1996), show that a Bohmian empty wave feels the interaction with an external gravitational field. This suggests the idea that, differently from classical mechanics, in which the particles are supposed to have an intrinsic mass, in Bohmian mechanics the mass is detached from the particle, entering as a dynamical parameter in the Schroedinger equation (according to the principle of parsimony). Put it succintly, the mass, in Bohm's theory, seems to be a physical property encoded in the dynamics of the wave function. Relevant attempts to clarify the situation can be found in Dürr et alii (2005), where the basic formalism of the theory is changed in order to account for a permutation invariant set of guidance equations and in Esfeld et alii (2014), where a philosophical explanation is given in terms of a non local holistic disposition grounded in the particles position of the total configuration of the Bohmian system.

The aim of the paper is to show how we can account for the emergence of a Newtonian classical particle from a Bohmian quantum one, i.e., how a massive particle might emerge at the macroscopic level starting from a massless particle at the quantum one.

The framework in background will be the Hamilton-Jacobi formulation of Bohm's theory. It permits to define the transition from the quantum non locality to the classical one in terms of the total potential: if the total potential depends only by the (inverse of the) distance between the particles, then we reach the classical non locality. Moreover, as working hypothesis, we will consider the PSI field as a real physical field, which guides the motion of the particles and gives rise to the correct probability distribution for the ensemble of the possible configurations of the system.

Within the general issue of the classical limit in Bohm's theory, we distinguish (at least) three different problems:

1. The emergence of the classical trajectories starting from the Bohmian ones.

2. The transition from the quantum non locality to the classical one.

3. The emergence of a massive classical particle from a massless Bohmian one.

Even though these problems are related each other, in the paper we mainly will focus on the third one.

From a dynamical perspective, the classical trajectories are obtained when we reduce the quantum modified Hamilton-Jacobi equation to the classical Hamilton-Jacobi one. For this part, we will mainly refer to the approaches, and some concrete results, by Holland (1993, section 6) and Bowman (2005).

Holland claims that the classical regime emerges when the quantum potential tends to zero and it is also slowly varying (so, also its gradient tends to zero). Following Bowman (2005), we can call these conditions the canonical conditions for the classical limit in Bohm's theory. Nevertheless, Bowman (2005) shows that the canonical conditions are neither necessary nor sufficient for the emergence of the classical dynamics, claiming, instead, that a more reliable condition is to describe the systems by narrow wave packets from the beginning, and introduce decoherence in order to keep the packets narrow in time. In our opinion, each of the two approaches have some virtues as like as some limitations. In particular, Holland's approach only considers closed and isolated systems, while we expect that a classical behavior emerges, in general, by the interaction with other quantum systems. On the other hand, Bowman's approach lies on the hardly justified assumption that the quantum systems, when we sought to account for the classical limit, are described by narrow wave packets from the beginning. Still, as before, if we consider that a macroscopic system is reasonably described by an open quantum system interacting with other quantum systems, then we are justified in considering narrow wave packets for the systems, because of the formation of mixed states due to the decoherence. In other words, decoherence comes first, explaining why we get narrow wave packets in a very short time. It is worth noting that, even if we think at the the wave function as a physical field in configuration space, in the classical limit, as a consequence of the action of the decoherence, the wave function of each system will be factorized. Under this circumstance, a projection of the PSI field in 3D physical space is quite straightforward.

Maintaining the classical regime of Bohm's theory as background, finally we will aim to clarify the ontological status of the mass in classical mechanics. For this purpose, we shall show that, for a macroscopic Bohmian system, the wave packet will be tightly picked around the center of mass of the system. The wave function does not disappear in the classical limit, simply it does not give rise to the characteristic quantum effects encoded by the quantum potential. This ontological picture follows from the universality and continuity of the quantum theory, which remains valid both for the microscopic and the macroscopic world.

At the classical level, the mass still remains a parameter encoded in the dynamical equation of the wave function, which is now the classical Hamilton-Jacobi equation. Since the wave function of a classical system is very narrow, the motion of the wave function and that one of the particle essentially overlap in 3D physical space, and we tend to identify the mass encoded in the wave function with an intrinsic property of the particle. Thus, Newtonian mechanics might emerge as an approximation of Bohm's theory, when the wave function of the system is very narrow (i.e., the coherence length of the packet is very small compared to the radius of the system) and tightly picked around its center of mass.

From a Bohmian perspective, the classical mass might be interpreted not as an intrinsic property of the particle, but as an emergent (dynamical) property of a more fundamental physical field.

Juha Saatsi and Lina Jansson. Varieties of abstract explanations: causal, non-causal, and mathematical

It is a commonplace that scientific explanations vary in their degree of abstractness: some explanations turn on very specific and concrete causal features, while others appeal to highly abstract and unspecific explanans. Yet others fall between the two extremes. But what does abstraction in a scientific explanation amount to? In this paper we will propose an answer to this question and relate it to two recently discussed issues in philosophy of explanation:

1. Are some scientific explanations *non-causal* by virtue of being sui generis 'abstract explanations'? (Pincock, BJPS, forthcoming)

2. Are some scientific explanations *mathematical* by virtue of being abstract in a suitable way? (Lange, BJPS, 2013)

One natural intuition is that an explanation's abstractness has to do with a lack of specificity: more abstract explanations have more abstract explanantia, which in turn can be (partially and comparatively) understood in terms of the 'number' (or measure) of possible cases to which the explanation applies (Weslake 2010). We will analyse this idea of abstraction-qua-lack-of-specificity in the context of a counterfactual account of explanation in the spirit (but not the letter) of Woodward (2003). We will argue that the counterfactual framework has natural conceptual resources for analysing abstraction in terms of the notion of 'same-object counterfactual': the more abstract the 'same object', the more abstract the explanation. We will cash out this thought by refining the crucial notion of 'same-object counterfactual'. In particular, we will pay attention to how a scientific theory fixes what counts as the 'same object' and identifies the relevant counterfactual in explanatory reasoning. This will allow us to incorporate broader theoretical considerations, including information about other objects, without abandoning the idea that what matters for explanatory purposes are the dependencies of the target system on the explanans.

As is well known, Woodward's counterfactual account of explanation (amongst others) regards as causal many highly abstract explanations (e.g. in biology and the social sciences). Therefore it is clear that one doesn't end up with a non-causal explanation on these views just by abstracting away from causal details. In the context of current literature an interesting question remains, however: are some explanations non-causal by virtue of being *sufficiently* abstract? (And if so, how abstract must such an explanation be?) Armed with our analysis of explanatory abstraction we argue that the answer the first question is 'no'. In this connection we will critically discuss Pincock's recent idea that some 'highly abstract' explanations – such as Konigsberg's bridges and Plateau's laws for soap bubbles – should be understood as *sui generis* abstract explanations. Part of Pincock's reasoning is that these are examples of explanations where the explanatory force stems from information about other relevantly similar systems. If this were the case then we would have a sharp distinction between these explanations and the causal ones in Woodward (2003). Contrary to Pincock, we will argue that a modified 'same-object counterfactual' analysis can naturally accommodate these very abstract explanations, so there is no reason to regard them as different in kind from other explanations that supply 'what-if-things-had-been-different' information. There are other reasons, however, for regarding these explanations as non-causal.

We will also briefly note how our analysis of explanatory abstraction relates to Lange's recent analysis of 'distinctly mathematical' scientific explanations. Lange (BJPS, 2013) presents a 'modal' account of mathematical explanations, according to which 'distinctly mathematical' explanations explain by reference to modal facts that are "modally more necessary than ordinary causal laws are." We contrast Lange's view with our counterfactual framework that views 'distinctly mathematical explanations' as abstract explanations that (i) locate the explanatorily relevant properties at an extremely abstract level, and (ii) use mathematics to represent the relevant 'same object'. We give various reasons to prefer our counterfactual analysis of the explanatoriness of such mathematical explanations over Lange's modal account.

Olivier Sartenaer. Emergent quasiparticles: the case of the fractional quantum Hall effect

Physicists often qualify quasiparticles as emergent entities or phenomena (Negele & Orland 1998). Among the different mechanisms or effects that can give rise to such putatively emergent quasiparticles, the fractional quantum Hall effect certainly occupies a prominent place, as one of its discoverer – the 1998 Nobel prize winner Robert B. Laughlin – suggests in his notorious and scathing pamphlet against physical reductionism: "If the quantum Hall effect raised the curtain on the age of emergence, then the fractional quantum Hall discovery was its opening movement" (Laughlin 2005, p. 76). Yet such strong claim hasn't received much attention in the community of philosophers of science, who usually discuss possible cases of emergence in physics in the contexts of chaotic systems (Newman 1996), non-relativistic quantum mechanics (Humphreys 1997), phase transitions (Butterfield 2011), string theory (Huggett & Würthrich 2013) or even – and somewhat surprisingly – classical mechanics (McGivern & Rueger 2010). Such an almost general omission of the fractional quantum Hall effect in the emergence/reduction literature might appear quite surprising, especially in the face of the fact that such an effect seems prima facie to gather all the features that any (even moderate) emergentist would want, viz. a (quantum) phase transition giving rise to a "new kind" of order (a so-called "topological" order), a "new" (topological) phase of matter, the advent of "new" (quasi)particles (anyons) that obey fractional statistics that are neither fermionic nor bosonic, "new" or "unprecedented" interactions (long-range entanglement) that could not have been deduced "from the bottom up", etc.

In the talk, we critically assess Laughlin's original claim that the fractional quantum Hall effect is a paradigmatic case – if not the case – of (a robust form of) emergence within physics itself. To this aim, we primarily reconstruct Laughlin's somewhat nebulous rhetoric in terms that are more adequate for a scrupulous philosophical discussion. In particular, we put some flesh on the bones of Laughlin's emergence ascription by explicating what kind of emergence he has – or should have – in mind, and this on the basis of a conceptual grid made of three cross-cutting distinctions, namely between epistemological and ontological, weak and strong, as well as synchronic and diachronic varieties of emergence. Through this procedure, we contend that the best way of capturing Laughlin's intuition about the fractional quantum Hall effect is through ontological, weak, diachronic emergence, a claim that sharply contrasts with Laughlin's (1999, 2005) own view that suggests to construe emergence in a synchronic, primarily holistic way – coherently with one of the rare existent philosophical exegesis of his work (Gillett 2010).

Ontological, weak, diachronic emergence actually constitutes one of the rare varieties of emergence for which no clear account and criterion are actually available on the philosophical market. Accordingly, we propose to build these on the basis of the particular case study that is the fractional quantum Hall effect, and this following the lines of Paul Humphreys recent – and still incomplete – account of emergence, called "transformational emergence" (Humphreys unpublished). As its name suggests, such an account construes the drive of any (putative) actual case of emergence involving a collective organization of some "parts" into a resulting "whole" as being an essential transformation of the properties of the parts, rather than a controversial, holistic and downward influence – be it causal or not – that the properties of the whole should exert on the properties of its parts, or a likewise controversial failure of mereological supervenience between the properties of the whole and the properties of its parts.

The contention that transformational emergence is well-suited to account for the striking features of the fractional quantum Hall effect is motivated by the actual physics involved in such an effect, and more particularly by a theoretical account of the effect that can be provided by the machinery of quantum field theory. In the lights of such a theoretical framework, the advent of all the "unexpected" or "unprecedented" features of the fractional quantum Hall effect is attributable to a suitable physical transformation (roughly, the confinement of an electron fluid) associated with a specific dimensional change (from 3+1 dimensions to 2+1 dimensions), whose effect is encoded in the presence of a new term – a Chern-Simons term – in the 2D-Lagrangian. We contend that the presence of such a new interaction term in the Lagrangian is a suitable criterion – a necessary and sufficient condition – for transformational emergence. A side advantage of this account is that it allows dissolving what appears to be a confusion in Laughlin's original rhetoric about the putative, principled non-deducibility of the features of the fractional quantum Hall effect from so-called "first principles". It turns out that there is some deduction available from "first principles", especially given the fact that it is possible in practice to theoretically account for the fractional quantum Hall effect in the lights of quantum field theory. But there is also a sense in which some "non-deducibility" is indeed at play, namely a non-deducibility between the physics of before the transformation (3+1 quantum field theory) and the physics of after the transformation (2+1 quantum field theory with the presence of a Chern-Simons term).

Construing emergence following these lines when it comes to an underappreciated phenomenon like the fractional quantum Hall effect should shed some new light on the notoriously complex relationship between particle physics and condensed matter physics.

References

Butterfield, J. (2011). Less is Different: Emergence and Reduction Reconciled. Foundations of Physics, 41(6), 1065-1135.

Gillett, C. (2010). On the Implications of Scientific Composition and Completeness. Or the Troubles, and Troubles, of Non-Reductive Physicalism. In T. O'connor & A. Corradini (eds.), Emergence in Science and Philosophy (pp. 25–45). New York: Routledge.

Huggett, N., & Würthrich, C. (2013). Emergent Spacetime and Empirical (In)coherence. Studies in History and Philosophy of Modern Physics, 44, 276-285.

Humphreys, P. W. (1997). How Properties Emerge. Philosophy of Science, 64(1), 1-17.

Humphreys, P. W. (unpublished). Transformational Emergence.

Laughlin, R. B. (1999). Nobel Lecture: Fractional Quantization. Reviews of Modern Physics, 71(4):863–874.

Laughlin, R. B. (2005). A Different Universe: Reinventing Physics from the Bottom Down. New York: Basic Books.

Negele, J.W., & Orland, C. (1998). Quantum Many-particle Systems. Boulder: Westview Press.

Newman, D. V. (1996). Emergence and Strange Attractors. Philosophy of Science, 63(2), 245-261.

McGivern, P., & Rueger, A. (2010). Emergence in Physics. In A. Corradini & T. O'Connor (Eds.), Emergence in Science and Philosophy (pp. 213-232). New York: Routledge.

David Schroeren. Theoretical equivalence as explanatory equivalence

Under what conditions are two theories equivalent? On its usual reading, this question is understood as an essentially semantic as opposed to a purely syntactic one: theoretical equivalence isn’t a matter of superficial similarities in the respective linguistic frameworks, but rather a matter of equality in content.

On an important strand of scholarship, the semantic dimension of theoretical equivalence is cashed out in intensional terms: given that certain logical or mathematical relations obtain between two theories, it follows that these theories are true at the same possible worlds. For instance, Clark Glymour has proposed that two theories are equivalent just in case they have a common definitional extension.

However, it seems that there is reason to be skeptical as to the sufficiency of such approaches for capturing our intuitive notion of theoretical equivalence. For one may expect that theories which are theoretically equivalent are also equivalent in explanatory contexts, in either of two ways. In epistemic sense of explanation, one might expect that the degree to which two equivalent physical theories exhibit cognitive virtues typically associated with successful scientific explanations - e.g. simplicity, parsimony, and unification - should be equal. On the other hand, in the metaphysical sense of explanation, one might expect that the sentences stating the reasons for why it is the case that, according to each of two equivalent theories, some (perhaps observational) fact obtains should be interchangeable. Put differently, if two theories are equivalent, we might expect them to posit, for each explanandum, the same set of worldly features in virtue of which that explanandum is explained by the explanans.

As has been recognised in the literature, explanation is not merely intensional, but hyperintensional. This means that necessarily or logically equivalent sentences cannot be substituted for the explanantia or explananda salva veritate, that is, without changing the truth value of the explanation. For instance, in the epistemic context of explanation, while the statement "That Phosphorus is in the sky is good reason to believe that it is morning" is true, the statement "That Hesperus is in the sky is good reason to believe that it is morning" is false, even though it is a necessary truth that Hesperus is Phosphorus. In the metaphysical context of explanation, the statement "3 is a prime number” is explained by the statement "What it is to be a prime number is to be divisible only by 1 and itself, and 3is divisible only by 1 and itself", the latter sentence is necessarily equivalent to (true at the same possible worlds as) all mathematical sentences. But substituting any mathematical sentence in the above explanation will generally make the explanation false: "3 is a prime number" is not explained by "1+1=2".

Now, this entails that, if we are right to expect that two equivalent theories should be equivalent in either of the epistemological or metaphysical contexts of explanation, then it is clear that none of the extant criteria for theoretical equivalence are sufficient; for any criterion of theoretical equivalence that is based only on merely intensional notions such as "...is interdefinable with..." or "...is equivalent in mathematical structure to..." must be blind to the hyperintensional nature of theoretical equivalence.

The purpose of this paper is to construct a formal criterion of explanatory equivalence on the basis of the hyperintensional context of scienti?c explanation. I shall be concerned specifically with the metaphysical notion of scientific explanation, intended to capture the way in which, according to some physical theory, worldly items and their features are causally or non-causally determined by other worldly items and their features.

The central thought underlying the construction of the proposed criterion is this. In order to understand the sense in which two scientific explanations of some fact are equivalent, we need to consider not merely the set of worlds at which some sentence of a theory is true, but rather the parts of the worlds in virtue of which that sentence is true at those worlds - its set of truthmakers or grounds, for short. This serves to introduce a criterion of explanatory equivalence, according to which two theories T and T' are equivalent (roughly,) iff, for each explanandum E, the explanans of E according to T has the same set of truthmakers as the explanans of E according to T'. I state this criterion formally in the framework of pure categorial languages, which allows me show exactly how explanatory equivalence goes beyond Clark Glymour's notion of definitional equivalence while incorporating the latter as a necessary condition. In this sense, the criterion of explanatory equivalence constitutes a refinement of Glymour's criterion, and for that reason, I suggest that the general thrust of the present proposal is best understood as a friendly amendment to the ongoing research projects on theoretical equivalence.

I begin by explicating the precise sense in which de?nitional equivalence is based on a merely intensional context. Subsequently, I show how the inclusion of hyperintensional two-place sentential functor 'because' in a pure categorial language results in a violation of semantic compositionality. Next, I define the notion of a hyperintensional valuation and use this to construct a natural criterion of explanatory equivalence of physical theories. Finally, I sketch an application of the criterion of explanatory equivalence.

Partly due to the above-mentioned weaknesses of the extant formal accounts of theoretical equivalence, it has recently been claimed that there is something wrong with these accounts qua formal. If the argument in this paper is successful, there is reason to think that this claim is premature.

Tobias Starzak. Morgan´s Canon - Interpretation and Justification

Questions about animal mentality are interesting to both philosophers and comparative psychologists and investigating the animal mind is thus an interdisciplinary enterprise. The genuine philosophical part is often thought to be the analysis of the relevant concepts involved – i.e. what it means to be conscious, think, possess concepts or have a theory of mind, to just mention a few of the mental capacities that play a role in this debate. However, we hardly find any unambiguous behavior and usually different interpretations concerning the underlying psychological processes are possible. Thus we need to look at the question of animal cognition from an epistemological and methodological point of view as well. The principle Morgan´s Canon, named after Conwy Lloyd Morgan, the "father of comparative psychology" (Karin-D´Arcy 2005, p. 179) fills this role:

"In no case is an animal activity to be interpreted in terms of higher psychological processes, if it can be fairly interpreted in terms of processes which stand lower in the scale of psychological evolution and development." (Morgan 1903, p. 59)

Morgan´s Canon might be the "most cited" (Dewsbury 1984, p. 187) and some even say the most important (Galef, 1996, p. 9) sentence in the history of comparative psychology and many philosophers as well as psychologists accept it (explicitly or tacitly) as a valid principle. It consists of two interconnected claims. First it claims that there is a scale of psychological evolution and development by reference to which we can classify different psychological processes as either higher or lower relative to each other. Secondly it claims that explanations that refer to lower psychological processes are to be preferred in cases where different interpretations are possible. In this talk I´ll argue that we can´t spell out the idea of a psychological scale (claim 1) in a way that satisfies the condition formulated in claim 2 and thus that Morgan´s Canon can´t be of any help in solving current debates concerning the animal mind.

In the first part I´ll focus on the idea of a psychological scale. Morgan´s own idea of such a scale is based on his understanding of evolution as a process that is always directed at increasing complexity and sophistication (Sober 1998). But this conception of higher and lower psychological mechanisms is undermined by the fact that it rests on an empirically inaccurate description of how evolution works. However, some philosophers defend the idea of a psychological scale on different grounds. One idea is to understand higher and lower psychological abilities as based on causal priority: a psychological process A is lower than a process B if A is a causal precondition for process B to evolve. Different versions of this idea have been defended (e.g. by Karin-D´Arcy (2005) and Shettleworth (2013)). A minor problem with this idea is that the processes it identifies as higher no longer necessarily correspond with more complex processes and thus that it changes Morgan´s Canon to a principle that´s very different from how it is usually understood. More importantly, however, it depends on the empirical claim that some processes are causally prior in this sense even across only distantly related species, which is questionable given the evidence we have for convergent evolution (Emery et al. 2004; Meketa 2014). Another suggestion from Hans-Johann Glock (forthcoming) suggests that we should understand psychological processes as higher or lower relative to each other on the basis of conceptual entailment relations. According to Glock a process A is higher than a process B if and only if having process A conceptually entails having process B. However, I´ll argue that this classification can´t be applied to many of the recent debates in comparative psychology, as for example the ToM debate, the debate about animal´s causal understanding or the question which processes of social learning underlie animal traditions and thus that it reduces Morgan´s Canon to a very restricted principle. Finally, on Elliott Sober´s account a psychological process A is higher than another process B "if and only if the behavioral output of process A properly includes the behavioral output of process B" (Sober 1998, p. 236). As the notions of a psychological scale discussed before, this idea offers a clear-cut criterion according to which we can plausibly rank psychological processes. But it faces similar problems, too. Adopting this definition of a psychological scale changes the meaning of Morgan´s Canon substantially. Moreover, as Sober argues himself, the question whether two psychological abilities in fact stand in the proper relation to each other is an open empirical question and it might turn out that this principle has similarly limited applicability as Glock´s proposal.

In section (2) I´ll examine the question whether a systematic preference of lower psychological abilities in explanations (for any of the interpretations discussed in Section (1)) can be justified. Most attempts to justify Morgan´s Canon do so with reference to parsimony or simplicity, or defend it as an attempt to circumvent a human tendency to anthropomorphize animal behavior (see Sober 1998; Karin-D´Arcy 2005; Fitzpatrick 2008). However, I´ll argue that none of these attempts is successful. Appeals to parsimony or simplicity face three interconnected problems. First, they can be understood in various ways and it is unclear how to apply these notions to Morgan´s Canon. Secondly, depending on which conception we apply there is no strong connection between the relative parsimony or simplicity of an explanation and it´s likeliness of being true. And thirdly, neither parsimony nor simplicity speak in all cases in favor of Morgan´s Canon and in some cases they even support a preference for explanations that refer to higher psychological processes. Finally, I´ll argue that using Morgan´s Canon to counter a human bias to anthropomorphize animals is irrational since it amounts to substituting one bias for another one.

If we dismiss Morgan´s Canon the question arises how we can decide between alternative explanations. In Section (3) I´ll suggest an outline of how we should deal with rival explanations solely on the basis of available evidence instead.

References

Dewsbury, D.A. (1984). Comparative psychology in the twentieth century. Stroudsburg, PA: Hutchinson Ross.

Emery, N.J. (2004). The Mentality of Crows: Convergent Evolution of Intelligence in Corvids and Apes. Science 306(5703), S. 1903–1907.

Fitzpatrick, S. (2008). Doing Away with Morgan's Canon. Mind & Language 23(2), S. 224–246.

Galef, B. (1996). Historical origins: the making of a science. In: Houck, L. & Drickamer, L. (Hrsg.). Foundations of Animal Behavior: Classic Papers with Commentaries. Chicago: Chicago University Press.

Glock, H.J. (2012). The anthropological difference: what can philosophers do to identify the differences between human and non-human animals. Royal Institut of Philosophy Supplement 70, S. 105 – 131.

Karen-D Arcy, M.R. (2005). The Modern Role of Morgans Canon in Comparative Psychology. International Journal of Comparative Psychology 18, S. 179–201.

Meketa, I. (2014). A critique of the principle of cognitive simplicity in comparative cognition. Biol. Philos. DOI 10.1007/s10539-014-9429-z

Morgan, C.L. (1903). Introduction to comparative psychology. London, UK: Walter Scott.

Shettleworth S.J. (2013). Fundamentals of comparative cognition. Oxford University Press, New York

Sober, E. (1998). Morgan's Canon. In: Cummins, D. & Allen, C. (Hrsg.). The Evolution of Mind. Oxford: Oxford University Press.

Thomas, R.K. (2001). Lloyd Morgan's Canon: A History of Misrepresentation. http://htpprints.yorku.ca/archive/00000017/00/MCWeb.htm#Author%20Note. Abgerufen am: 15.10.2013.

Gordon Steenbergen. Cognitive Neuroscience and the Mechanist Thesis

Cognitive neuroscience is an interdisciplinary enterprise aimed at explaining cognition and cognitive behavior. It appears to be succeeding. What accounts for its apparent explanatory success? One prominent philosophical thesis is that cognitive neuroscience explains by discovering and describing mechanisms. In this essay, I identify and critically assess the theoretical commitments of an important interpretation of this thesis that Carl Craver and David Kaplan defend. On their view, the mechanist thesis is defensible on both descriptive and normative grounds: cognitive neuroscience is in the business of describing mechanisms; and mechanistic descriptions, insofar as they describe the network of causal dependencies that produce a cognitive phenomenon, are paradigm examples of good explanations. Indeed, they go so far as to argue that mechanistic descriptions are necessary for explaining cognitive phenomena. However, I claim that arguments in defense of these commitments fall short of their normative and descriptive aims. Normative arguments for the necessity of mechanistic explanations are not persuasive: even supposing that mechanistic descriptions are required to adjudicate among explanatory hypotheses in the mind sciences (a controversial assumption), it does not thereby follow that the model that does the explaining must itself be mechanistic. Furthermore, the explanatory variety that is characteristic of cognitive neuroscience poses a significant challenge to the descriptive claims of the theory as Kaplan and Craver understand it. I support these claims by considering research on the neuroscience of decision-making as a representative example of research in this field. Instead, an alternative account of the role of mechanisms in explaining cognitive phenomena emerges, namely, that mechanisms are a rich source of evidence that can be marshaled in support of a variety of kinds of explanatory models.

Chris Timpson. Bell's theorem, local causality, explanation and Everett

Last year – 2014 – was the 50th anniversary of Bell's theorem. In those 50 years, a very great deal of high-quality work has been done investigating the structure, implications, and various extensions of the theorem. But something surprising emerged from the stimulating discussions surrounding the 50th anniversary: the extent to which - even now - there is still no consensus amongst experts about what Bell's theorem actually shows. (See, e.g. J Phys A Special Issue, 50 Years of Bell's Theorem vol 47, no. 42 (2014); Int J Q Found, John Bell Online workshop 2014 http://www.ijqf.org/groups-2/bells-theorem/forum/)

One of the main points still under vigorous dispute is whether - or in what sense - violation of Bell inequalities shows a theory, or shows the world, to be nonlocal.

One view, trenchantly expressed e.g. by Maudlin and Norsen (see refs. above), maintains that Bell's condition of local causality is the one correct formulation of the notion of locality, and that (modulo the standard logical loopholes, such as the mathematically equivalent trio of superdeterminism, retrocausation or non-common-common cause) violation of a Bell inequality by a theory simply shows that that theory is nonlocal; and experimental violation of the inequality simply shows that the world is nonlocal.

Others demur, perhaps pointing out that the notion of locality ought to be divided into a number of distinct ideas, e.g., separability vs local-action (Howard, Healey and others); and it may not be clear which of these's failure is implicated by failure of local causality (thus, by violation of a Bell inequality); whilst, furthermore, it may remain unclear whether any such failures, if failures there be, would amount to a *worrisome* form of nonlocality – whether from the point of view of quotidian metaphysics or from the point of view of relativistic covariance. Or again, it might be asserted that the conclusion to the presence of a worrisome form of nonlocality from violation of a Bell inequality only holds when certain kinds of theories are considered – e.g. deterministic hidden variable theories (a la de Broglie—Bohm), of the kind Bell originally considered in 1964.

Maudlin (2010) diagnoses this last kind of view – the view that only theories of a certain kind, kind X, are shown to be nonlocal if they violate a Bell inequality – as being subject to the 'fallacy of the unnecessary adjective'. As he sees it, the qualification 'X' in 'theory of kind X' is superfluous, for any theory *at all* which violates a Bell inequality is simply going to be nonlocal. But others remain unpersuaded by this contention, not least as there appear to be concrete counterexamples to the claim. That is, there seem to exist examples of theories which violate Bell inequalities yet which are not nonlocal, or which are not nonlocal in any worrisome sense.

One such class of theories are operationalist or instrumentalist in flavour, denying descriptive content to much of the formalism of quantum theory, and resting with the no-signalling theorem (which automatically holds in quantum mechanics) as the only interesting, or the only acceptable, statement of locality for the theory. But more striking to the realist minded is the example given by Everettian quantum mechanics.

Certainly, it is a widespread view amongst those who have entertained the Everett interpretation seriously that Everettian quantum mechanics does not involve any nonlocality in the dynamical sense - the sense of action-at-a-distance. (Cf., Everett, 1957; Vaidman 1994; Bacciagaluppi 2002; Timpson and Brown 2002; Wallace 2012; Tipler 2014; Brown and Timpson 2014.) Granted, the theory is kinematically nonlocal – it is non-separable – but whilst striking, this feature is not pathological, nor physically worrisome, especially since it sits entirely straightforwardly within a Lorentz-covariant setting.

In this paper, I shall review the dialectic to date, and defend the claims of the Everett interpretation to provide an account of the world which violates local causality (violates Bell inequalities) whilst avoiding any form of dynamical nonlocality, that is: action-at-a-distance. In short, I shall defend the claim, contra Maudlin and Norsen, that there exist interesting examples (at least one!) in which failure of local causality does not entail the presence of nonlocal causes, or action-at-a-distance.

As I shall seek to elaborate, two features intertwine in the Everettian, local, story – the absence of definite outcomes for a far measurement in an EPRB-Bell experiment relative to definite outcomes at the near side (i.e., the presence of multiple superposed outcomes, from the God's eye view), and the fact that the theory is non-separable (in its incorporating entanglement). However, Henson (2013, 2014) has provided pungent arguments that neither of these features can, generically, save a theory from being (dynamically) nonlocal when it violates a Bell inequality. I shall seek to rebut his worthy contentions by, first, offering a reasonable definition of 'local' which fits a non-separable context, and second, by drawing-out how his argument has been misled by too narrow a formalisation of the important intuition captured by Reichenbach's principle of the common cause.

In his final writings, Bell (1990) was much concerned by the thought that if Bell-inequality violating correlations could after-all be deemed local, if that is, his notion of local causality did not in fact capture the appropriate sense of physical locality, then one would be left with the unlovely prospect that there could be robust statistical correlations predicted within one's theory (and, indeed, found in the world) for which no explanation could be found. For local causality was supposed to be the mathematical formalisation of the very plausible idea that the proximate causes of events ought to be nearby them, and causal chains leading up to events should be confined (in a relativistic setting) within their past light cones. But correlations which violate a Bell inequality cannot be explained by such causal chains within the past light-cones of the measurement events, and if – even so – they are local, then they cannot be explained by any (dynamically) nonlocal influence either. It would seem to follow that they cannot be explained *at all*. Some (e.g. Fine, van Fraassen) have embraced this conclusion, but Bell found it repugnant, and I would concur. Happily, a trick has been missed in Bell's reasoning. In the context of theories which may be non-separable, then one must not implement the intuitive Reichenbachian causal thought (viz., correlations should be explicable either by direct or by mediate causal influence – common cause) in the way that Bell (and perhaps Reichenbach) presumed. Dynamically evolving non-separable global states can provide the required explanation of correlation, even when no (traditional) Reichenbachian common cause, and no instantaneous proximate cause (action-at-a-distance) can be found. In Everett one finds precisely an instance of this more general structure.

Nick Tosh. Ensemble realism: A new approach to statistical mechanical probability

"What we know about a body can generally be described most accurately and most simply by saying that it is one taken at random from a great number (ensemble) of bodies which are completely described."

So wrote Josiah Willard Gibbs in 1902, but he didn't quite mean it. Ensembles, for Gibbs, are convenient fictions: the mug of tea on your desk doesn't really have trillions of doppelgangers, but for certain purposes——defining and calculating statistical mechanical probabilities——it is useful to imagine that it does. Gibbs's formal apparatus survived the twentieth century unscathed. It remains the go-to tool for handling statistical mechanical probabilities, and these probabilities play ineliminable (though not always acknowledged) roles in practically all scientific inferences. Nevertheless, no one has yet been tempted to ditch Gibbs's antirealism and take talk of ensembles seriously. Why not? After all, the foundations of statistical mechanics are vigorously disputed; the literature is large and diverse; and philosophers of science are not usually shy about pushing (or at least entertaining) bold realist hypotheses. I can think of five reasons for reluctance in this case.

(1) No one has ever seen an ensemble.

(2) It is not clear how to make sense of the suggestion that bodies are 'taken at random from' ensembles.

(3) Ensemble realism would lumber us with a naïve actual-frequentist account of statistical mechanical probability; actual frequentism is widely assumed to be false (see, e.g., Hájek [1997]).

(4) We expect ontologically extravagant hypotheses to be motivated by 'spooky' phenomena (the Everett interpretation seems outrageous until we read about the double-slit experiment). But the phenomena explained by statistical mechanics——cooling tea, melting ice cubes——are utterly mundane.

(5) There is lingering sympathy for the view that statistical mechanical probabilities are in some sense subjective or psychological.

These reasons, I contend, do not add up to a vindication of antirealism. (4) and (5) can be dismissed quickly. The phenomena explained by statistical mechanics are mundane only in the sense that they are familiar; familiar phenomena sometimes have surprising explanations. And it is an objective, non-psychological fact that hot tea tends to cool. If statistical mechanical probabilities help explain this fact, they had better be objective. That leaves (1)–(3).

I have argued elsewhere that actual frequentism does not deserve its poor reputation. Specifically, in (Tosh [forthcoming]) I defend an actual-frequentist analysis of non-deterministic chance (the kind associated with quantum mechanics). I now wish to adapt that analysis to cover statistical mechanical probability. The result will——I hope——be a version of ensemble realism that can deal with (1) and (2) above in a principled way. In the remainder of this abstract, I shall offer brief sketches of the original analysis, the adaptation strategy, and my grounds for optimism.

According to my analysis of non-deterministic chance, the existence of non-trivial chances implies the existence of causal pasts (roughly, boundaries and interiors of past light cones) that are perfect intrinsic duplicates. I identify chances with relative frequencies defined over classes of such duplicates. So, for example, if Fred's present chance of living to be 100 is 0.31, then the pattern of events in Fred's causal past is instantiated finitely many times, and the man playing the Fred role lives to be 100 in 31% of these cases. The main selling point of the analysis is that it allows us to ground the Principal Principle in self-location indifference. Admissibility conditions work out as one would hope: for example, historical hypotheses are admissible, because Fred and his counterparts have matching histories. The analysis faces the obvious objection that duplicate casual pasts——which would be large regions of space-time——have not been detected by astronomers. However, it turns out that we should not expect to observe duplicate pasts, even if many exist. The reason is geometric: light-cone interiors that are intrinsic duplicates cannot overlap unless they have peculiar internal symmetries. Generically, then, such regions will be disjoint.

To obtain an analysis of statistical mechanical probability, only one modification is necessary: we must coarse-grain the equivalence relation that generates the reference classes. The resulting analysis identifies statistical mechanical probabilities with relative frequencies defined over classes of macroscopically equivalent causal pasts, where two causal pasts are macroscopically equivalent iff they are intrinsic duplicates modulo coarse-graining. (Such frequencies depend not only on the system of interest, but also on the coarse-graining scheme. Since statistical mechanical probabilities are relative to levels of description, this is a feature rather than a bug.) The geometric explanation for our failure to observe duplicate causal pasts remains compelling, unless the coarse-graining is very coarse indeed. Moreover, self-location indifference will continue to ground Principal-Principle-shaped credence constraints, so long as epistemic agents are unable to discriminate between macroscopically equivalent histories. Admissibility conditions then work out as expected: macro-level hypotheses are admissible; micro-level hypotheses are not. It will be rational for an agent to regard herself (and hence also whatever system she is studying) as 'taken at random' from an ensemble of macrohistoricallyequivalent duplicates in precisely those cases where we would expect statistical mechanical probabilities to be authoritative vis-à-vis agents' credences.

If ensemble frequencies are statistical mechanical probabilities, then they had better be moderately robust with respect to coarsenings, refinements and other changes in coarse-graining schemes. The analysis I have sketched does not guarantee this. But nor should it! Robustness is, at best, nomically necessary; metaphysically it is contingent. Laws governing statistical mechanical probabilities would, if my proposal is correct, be global constraints on ensemble frequencies: 'spooky' perhaps, but——thanks to the generic remoteness of duplicates——not observably so.

References

Hájek, A. [1997]: '"Mises Redux"-Redux: Fifteen Arguments Against Finite Frequentism', Erkenntniss, 45, pp. 209-27.

Tosh, N. [forthcoming]: 'Finite Frequentism in a Big World', British Journal for the Philosophy of Science (http://bjps.oxfordjournals.org/content/early/2014/12/11/bjps.axu027).

Dana Tulodziecki. From zymes to germs: discarding the realist/anti-realist framework

According to one of the main anti-realist arguments, the pessimistic meta-induction, we have reason to believe that our current theories are just as false as their predecessors. Proponents of this argument draw attention to a list of theories that were once regarded as highly successful, yet ended up being discarded and replaced by radically different ones. Scientific realists, in response, have argued, first, that the anti-realists' list is too permissive, and ought to be restricted only to theories that enjoyed `genuine' success, which, according to realists, consists in a theory's ability to make (use-) novel predictions, i.e. predictions that played no role in the generation of the original theory. Second, in dealing with the remainder of the so diminished list, realists have proposed and endorsed a variety of selective realisms (notably those of Kitcher, Worrall, and Psillos) which emphasise the carrying over of stable and continuous elements from earlier to later theories and which are then used to argue for the approximate truth of those earlier theories.

In this paper, I argue that neither realist nor anti-realist accounts of theory-change can account for the transition from zymotic views of disease to germ views. I begin by explaining the zymotic theory of disease, one of the most sophisticated and popular versions of the mid-nineteenth miasma theory. The zymotic theory theory drew on some of the most successful science at the time, such as Liebig's chemical theories, thereby allowing it to propose highly detailed mechanisms about the exact manner of disease causation. According to the zymotic theory, diseases occur as a result of introducing into the body various zymotic materials, either through direct inoculation or through inhalation after being dispensed in the air. Essentially, these zymotic materials were thought to be putrefying organic matter that would communicate its process of decomposition to pre-existing materials in the victim's blood where it would act in a manner similar to ferment, thus causing diseases.

After explaining the basics of the zymotic theory, I then show (i) that the zymotic theory and it successor, the germ theory, are strikingly different in almost every respect and (ii) that, despite the fact that the zymotic theory was so different from its successor, it was highly successful. Moreover, I show (iii) that this is so even according the realists' own, more stringent, criterion of success as consisting of use-novel predictions. Some examples of such use-novel predictions were the zymotic theory's predictions about what geographical regions ought to be affected by diseases to what degrees, and, strikingly, a number of numerically very precise predictions resulting from Farr's so-called elevation law of 1852, relating cholera mortality and the elevation of the soil. Other novel predictions concerned the course and duration of epidemics, the relation between population density and disease morbidity and mortality, the relation between mortality rates and different occupations, and relations between mortality from various diseases and age.

I argue, however, that despite the zymotic theory's successes, realists cannot account for the zymotic case. According to selective realists, precisely those parts that were indispensable to a theory's genuine success are the ones that ought to be retained; yet, as I show, there is no discernible continuity between the zymotic theory and the germ theory: the zymotic theory had an entirely different ontology and structure from that of the germ theory, and it was also radically conceptually different in other ways, such as in its focus on processes of decay as opposed to pathogenic entities. Thus, there were no stable or invariant elements that were carried over from the zymotic to the germ theory: neither its entities, nor its mechanisms or laws, nor its processes, or even the structure of diseases themselves was retained.

It thus appears that the zymotic theory is exactly the kind of case that anti-realists are looking for as support for the pessimistic meta-induction: it was highly successful, discarded, and had very little in common with its successor. However, I argue that, in fact, anti-realists fare no better than realists, since there was also no radical conceptual change or discontinuity between zyme and germ views: despite the fact that the zymotic theory and the germ theory -- viewed as finished products -- are radically different, the transition from the former to the latter was neither radical nor sudden.

To make this point, I show that there were no clearly defined and opposing germ and anti-germ research programmes, as is often claimed; in particular, there was no switch from one of these views to the other, but, instead, a gradual transition during which different aspects of a number of germ views were slowly assimilated into the zymotic theory. Elements of zymotic and germ views co-existed for some time, until, eventually, various parts of the zymotic theory were discarded, little by little, as increasingly well-defined versions of the germ theory emerged and started taking hold. The specific examples I use to argue for this position are (i) the changing views about the media of disease transmission, (ii) the changing views about the nature of zymes, and (iii) the change from chemical views of disease to biological ones.

I conclude that neither realist nor anti-realist views can adequately account for the transition from zymes to germs. However, I argue that the problem lies not with specific realist or anti-realist proposals, but, rather, with an unwarranted assumption they both share, namely the assumption that there are well-delineated theories that can be compared and assessed on terms set by the realism-debate in the first place, an assumption that does not hold in this case.

Peter Vickers. No Miracles? Scientific Realism and the 1811 Gill Slit Prediction

In the years since the heyday of the 'no miracles argument' (Putnam, Boyd, Leplin, etc), scientific realists have had to re-think what would and would not count as a 'miracle'. The original argument asked us to accept that if scientific theories were not (basically) true, they could not possibly enjoy the success that they do (that would be a 'miracle'). Today the locution 'no miracles argument' is somewhat out of favour, but the argument remains within the realist literature. Now the (often implicit) claim is something more like the following: a scientific theory could not possibly enjoy *novel predictive* success if the hypotheses *doing work* to generate that success were not at least approximately true (that would be a 'miracle'). The realist was forced to adjust her position in response to episodes in the history of science where we (apparently) find exactly what was supposed to be sufficiently unlikely to be described as 'miraculous'. However, historical case studies are still forthcoming which (apparently) challenge even the contemporary realist positions.

One such case concerns J. F. Meckel's 1811 prediction of gill slits in the development of the human embryo. To see whether this case is relevant to the debate our first question must be whether this prediction is truly a *novel predictive* success, or whether Meckel somehow devised his theory specifically to reach this prediction. In fact, the prediction was temporally novel: the phenomenon was completely unknown at the time Meckel made his prediction. In 1825 Heinrich Rathke discovered gill slits in pig and chick embryos, and eventually the gill slits in the human embryo were observed in 1827 by Rathke, von Baer, and others.

Having established that the prediction is novel in the sense required, the next question is whether the 'working posits' involved in reaching the prediction are (approximately) true. Our answer starts with the fact that Meckel was part of the Naturphilosophie movement within Germany at the turn of the 19th century. This movement was based on some basic assumptions concerning biological development:

(i) The animal series assumption: Animals can be put in a series from 'lower' to 'higher', with *human being* at the top of the pile.

(ii) Teleology: 'Lower' animals are striving to be human beings, and fail to become human beings because their development halts (Oken described lower animals as 'human abortions').

(iii) A 'single biological force' assumption: A single force or 'power' underlies all biological development. Thus similar organisms must develop in the same general biological 'direction' (remaining similar after a period of development).

(iv) A developmental theory of the animal series: The series of animals from 'lower' to 'higher' is caused by each animal developing from the same starting place, but to a different extent.

From here it is a very short step to the conclusion that the series of adult organisms from 'lower' to 'higher' must parallel the stages of human development from a 'primal zero' or 'initial chaos' (Oken) to a final, adult human being. Thus there is a period where the human embryo is a fish, and it follows that there is a period where the human embryo has gill slits.

Can the realist claim that the hypotheses doing work here are (approximately) true? Prima facie it seems highly unlikely, simply because there is so little within Meckel's assumptions that is today considered (approximately) true. It goes radically against evolutionary biology to try to place animals in a single series from 'lower' to 'higher', and teleology and the 'single biological force' assumption have no place in modern developmental biology. Even Meckel's parallelism is misguided: it is *not* the case that the development of the human embryo parallels (even roughly) the (alleged) animal series. Thus one might describe Meckel's predictive success as 'lucky'. And, if the realist thinks it's a miracle to get lucky in this way, then there *are* miracles in science!

How can the realist respond? As before in this debate, the realist might adjust her position, perhaps insisting that for realist commitment we need *quantitative* novel predictive success. The problem here is that it starts to look like the realist position is too flexible to ever really be challenged, thus making nonsense of the whole debate.

But a better realist response might just be possible. It turns out that Meckel's 1811 prediction is phrased as follows: "Perhaps [Vielleicht] there is even a much earlier period when the embryos of the higher animals are also furnished with inner gills." With the use of the word 'vielleicht' it is clear that Meckel is speculating here, as opposed to making the sort of deductive prediction we are familiar with in the realism debate. To put it another way, the prediction isn't 'risky' in the way the prediction of the Poisson white spot was for Fresnel's theory of light.

This reduces the significance of the prediction somewhat, but it is still telling that Meckel was confident *enough* to state this (tentative) prediction in print. However, the realist can add to this consideration the suggestion that, if we more fully contextualise Meckel's prediction, we can come to understand why it wasn't so surprising that he reached this idea. The realist strategy here is to argue that, given the purely empirical/observational knowledge of the day, one could – with a bit of imagination – come to the gill slit prediction without *any* substantial theoretical ideas. For example, Meckel knew of a stage of development of the human heart where it very closely resembles a fish heart. And he also knew that frogs – which breathe with lungs – have gills at an early stage of their development (when they are tadpoles). Thus perhaps the realist can argue that Meckel really reached his conclusion not via his (false) theoretical ideas, but rather via his empirical knowledge.

This paper investigates these issues in detail, considering in particular whether the realist has a convincing response here, and, if so, whether this constitutes an(other) adjustment to the contemporary realist position, and the underlying ‘miracle’ intuition.

Ioannis Votsis. How to Make a Long Theory Short: Lessons from Confirmation

Scientists tend to opt for simpler and more unified theories. In this talk, I put forth a novel conception of unification as well as an associated formal measure. I begin the discussion with a brief survey of some failed attempts to conceptualise unification. I then proceed to offer an analysis of the notions of confirmational connectedness and disconnectedness. These are essential to the proposed conception of unification. Roughly speaking, the notions attempt to capture the way support flows or fails to flow between the content parts of a theory. The more the content of a theory is confirmationally connected, the more that content is unified. Theories that make more strides toward unification, and, hence, are more economical in the way they capture the same phenomena, are thus to be preferred to those that make less strides for purely confirmational reasons.

Attempts to devise a satisfactory conception of unification abound. One of the earliest is Friedman (1974) where it is argued that understanding is generated when we reduce the number of independently acceptable law-like assumptions that feature as explanantia in the derivation of an explanandum. The lower that number the more unified an explanation. Friedman’s account was in great part motivated by a desire to avoid trivial explanations. It had already been observed that deriving an explanandum from a set of premises is not sufficient to turn those premises into a genuine explanation. Friedman sought to avoid this problem by limiting the derivations that yield genuine explanations to those that unify phenomena. Though highly influential, his account soon faced a number of insurmountable difficulties. As Kitcher (1976) and others pointed out, Friedman’s account rules out trivial explanations only at the expense of also ruling out some genuine ones. Several other attempts at conceptualising unification have been made with similar problems. They include Forster (1988), Kitcher (1989), Schurz and Lambert (1994) and Thagard (1993).

While it ultimately fails, Friedman’s account does at least get one fundamental thing right. By emphasising the role of the acceptability of law-like assumptions his account places a premium on the link between unification and confirmation. The proposal in this talk agrees with this appraisal and indeed elevates the link with confirmation to the single most important ingredient in our quest to understand unification. According to this account, unification is to be understood as a measure of confirmational connectedness. But what is confirmational connectedness and its opposite confirmational disconnectedness?

Roughly speaking, the notions attempt to capture the way support flows or fails to flow between the content parts of a theory. The more the content of a theory is confirmationally connected, i.e. support flows between its content parts, the more that content is unified. Let us use ‘$x \vdash_r y$’ to denote that y is a relevant deductive consequence of x. In formal terms, confirmational connectedness can be articulated thus:

Any two content parts of a non-self-contradictory proposition G expressed as propositions A, B are confirmationally connected if, and only if, for some pair of internally and externally non-superfluous propositions a, b where $A \vdash_r a$ and $B \vdash_r b$: either (1) where $0 < P(a), P(b) < 1$, $P(a/b) \neq P(b)$ or (2) there is at least one true or partly true atomic proposition c such that $a \& b \vdash_r c$, $a \nvdash_r c$ and $b \nvdash_r c$.

An explication of the notions in the analysandum cannot be pursued in the abstract due to obvious limitations of space. Suffice it to say that the probabilities are meant to be objective. That is, probability statements indicate true relative frequencies and/or true propensities of things happening like events, states-of-affairs or property instantiations. An objective interpretation of the probabilities captures the intuition that the confirmational (dis-/)connectedness of the content of a theory is determined by facts about the world, i.e. it is not a subjective matter.

We are now ready to express the unification u of a proposition D with the following function:

u(D) = 1 - SUM(d_i(a,b)|i=1 to n) / SUM(t_i(a,b)|i=1 to n)

where d_i(a,b) denotes the number of disconnected pairs a, b in a given content distribution i, t_i(a,b) denotes the total number of connected plus disconnected pairs a, b in a given distribution i and n denotes the total number of content distributions. To determine the number of disconnected pairs in a given content distribution we count how many times a different pair of relevant deductive consequences a, b fails to satisfy either clause (1) or (2). Any pair that is not disconnected is counted as connected. The higher the value of u(D) the more unified its content. That's how you (justifiably) turn a long theory short, i.e. by insisting that it's content is confirmationally connected.

References

Forster, M. (1988) 'Forster Unification, Explanation and the Composition of Causes in Newtonian Mechanics', Studies in History and Philosophy of Science, vol. 19: 55-101.

Friedman, M. (1974) Explanation and Scientific Understanding. Journal of Philosophy 71(1): 5-19.

Kitcher, P. (1976) Explanation, Conjunction, and Unification. Journal of Philosophy 73(8): 207-212.

Kitcher, P. (1989) 'Explanatory Unification and the Causal Structure of the World' in Scientific Explanation: Minnesota Studies in the Philosophy of Science, P. Kitcher and W. Salmon (eds.), Minneapolis: University of Minnesota Press, pp. 410–505.

Schurz, G. and K. Lambert (1994) 'Outline of a Theory of Scientific Understanding', Synthese, vol. 101: 65-120.

Thagard, P. (1993) Computational Philosophy of Science, Cambridge, MA: MIT Press.

Gregory Wheeler and Conor Mayo-Wilson. Epistemic Decision Theory's Reckoning

Epistemic decision theory (*edt*) is a reform movement within Bayesian epistemology that aims to provide \emph{purely epistemic} criteria for evaluating the rationality of our beliefs, much in the same way that traditional decision theory provides criteria for evaluating the pragmatic rationality of our actions.

The driving force behind *edt* is a sharp distinction between the practical rationality of decision-making and the epistemic rationality of belief. All things considered, my *decision* to drink tea rather than coffee in the morning is rational just in case I prefer tea to coffee. But my preference for tea is neither here nor there when assessing the accuracy of my *belief* that I will drink tea in the morning. The accuracy of that belief depends instead on whether there is tea in the pantry, not upon my preference. So a fundamental tenet of *edt* is the conviction that epistemic value is distinct from any particular individual's subjective preferences. We call this ascetic devotion to epistemic purity *epistemic puritanism*.

Although the aims and justification of *edt* differ from that of traditional decision theory, both share the same mathematical heritage. Both typically represent agents' beliefs by probability functions, for instance. Further, just as traditionalists argue that there is a numerical function quantifying the utility of each action in each world, so epistemic decision theorists maintain that there is a numerical function quantifying the “epistemic utility" of a belief state in any given world. Thus, a second tenet of *edt* is that epistemic value is *numerically quantifiable*.

The problem for *edt* is that epistemic puritanism and quantifiability are incompatible. Our thesis is that epistemic puritanism subverts the assumptions necessary to establish a numerical representation of epistemic utility because those assumptions depend on features of a rational agent's subjective preferences. Our argument reveals the underappreciated strength of “pragmatic" arguments for probabilism, for it is precisely their assumptions about preference that make quantifying utility remotely defensible.

Our longer paper is structured as follows. Section one reviews the motivation for epistemic puritanism. In section two, we summarize the assumptions necessary for representing preference via a numerical utility function and review how those assumptions are defended in traditional decision theory. In particular, we focus on four core axioms of von Neumann and Morgenstern's \citeyear{NM:1944} governing preference: \textbf{totality}, \textbf{transitivity}, \textbf{independence}, and \textbf{continuity}---the last of which itself includes a structural axiom on the options of choice, namely that the set of options for choice include \textbf{lotteries}. In section three, we present our argument that puritanism is incompatible with the claim that epistemic utility is numerically representable. In short, \f{edt} faces a dilemma. Either the arguments for an epistemic utility function go through by slipping the constraints of epistemic puritanism, in which case \f{edt} is empty; or the arguments for an epistemic utility function fail because epistemic puritanism undercuts the means for justifying the axioms of the representation theorems. Finally, we consider and reject puritanism without quantifiability.

Alastair Wilson. Naturalizing Recombination

Various realist theories of the nature of modality appeal to some version of a principle of recombination. The informal ‘patchwork principle’ of Lewis’ On the Plurality of Worlds and the more rigorous treatment in Armstrong’s A Combinatorial Theory of Possibility are familiar examples, but recombination principles feature in most views according to which possible worlds are real structured entities. However, surprisingly little attention is usually paid to the epistemic status of these principles. They are paradigm mysterious examples of the putative synthetic a priori - highly substantive truths about the nature of modal reality, our way to knowledge of which we are somehow supposed to be able to reason. Now that transcendental idealism is out of fashion, and conventionalism has had its day, the most plausible treatment of the epistemic status of putative synthetic a priori truths is the Quinean one: such truths are justified through their indispensable contributions to the formulation of our best scientific theories. Nevertheless, the aura of mystery remains: the justification for such truths seems to be different in kind and much less direct than the justifications we have for truths about goings-on in the actual world.

The best way to render putative synthetic a priori truths unmysterious is to naturalize them. This means finding for them a home within the scientific worldview, rather than merely showing that they are necessary to underpin that worldview. In this paper, I offer a naturalistic treatment of recombination in the context of the Everettian, or many-worlds, approach to quantum mechanics. According to my proposal, the unitary evolution described by the Schrödinger equation is best understood as more akin to a recombination principle than to a law of any individual world. The Schrödinger equation grounds the truth of a macroscopic recombination principle, and thereby ensures that we have what David Lewis has called a plenitude of possibilities.

The Schrödinger equation has always had a puzzling status; it does not fit neatly into familiar frameworks for understanding laws of nature. Taken literally, it seems to describe an entity - the wavefunction - which evolves in an unfamiliar infinitely-high-dimensional Hilbert space. Even if the wavefunction is interpreted in some more palatable way - for example, Bohmians often think of it as something like a law - then it still encodes a remarkable amount of complexity: the wave-function of the universe, even for Bohmians, encodes enough information to reconstruct the entire space of physical possibilities (at least, those with the same initial conditions). If the Schrödinger equation is a law of nature, it is a law unlike any other.

Naturalizing recombination involves treating Everett worlds as distinct possible worlds: many-worlds quantum theory is then a theory of modality rather than merely a theory of the actual world. Unlike other versions of modal realism, Everettian modal realists have available a fully naturalistic story about the extent and contents of modal reality. This opens the way to a previously unsuspected possibility in metametaphysics: we can have empirical confirmation and disconfirmation for theories of the metaphysics of modality.

The reading of the Schrödinger equation as naturalizing recombination can in fact be motivated from within David Lewis' own view of the correct methodology for theory-building in science, as encoded in his 'best-systems' approach to laws of nature. If (as committed modal realists should be ready to) we expand the scope of this methodology to cover theories about the nature and contents of modal reality, then general facts about the space of worlds that strike good combinations of simplicity and strength will be good candidates to be included in the laws of nature. This conclusion is not limited to Everettians: modal realists in general can think of their preferred principle of recombination as a law of the plurality.

Lewis appealed to the principle of recombination in order to capture what he called the requirement of plenitude: a modal realist should ensure that there are "worlds enough, and no gaps in logical space". The Schrödinger equation plays a very similar theoretical role in Everettian quantum mechanics: it can be thought of as a conservation principle for probability. It ensures that there is an Everett world for every outcome that - before the experiment - had non-zero probability, and it thereby ensures something very like Lewis' plenitude of possibilities. Still, the theoretical roles are not identical: whereas the Lewisian patchwork principle applies to fundamental entities, the plenitude of possibilities guaranteed by the Schrödinger equation is at the derivative level.

The best modern versions of Everettian quantum mechanics lean heavily on aspects of the physics of decoherence. According to decoherence-based versions of the interpretation - associated in particular with Simon Saunders, David Wallace and Hilary Greaves - macroscopic worlds are derivative entities, somewhat indeterminate in their nature and number. Accordingly, Everettian modal realists ought to regard the plenitude of qualitative possibilities at the macroscopic level as a corollary of a more basic principle: fundamental reality evolves unitarily. Here, then, we have a putative truth of physics which is clearly more fundamental than a putative truth of metaphysics. The rationalist picture of metaphysics underlying physics, promoted by metaphysicians like George Bealer and E.J. Lowe, is fully inverted.

Nicolas Wüthrich. Conceptualizing uncertainty: An assessment of the latest uncertainty framework of the Intergovernmental Panel on Climate Change

Climate change has the potential to generate tremendous ecological, economic, and social impact. Although it is commonly agreed that this change is also driven by anthropogenic factors, it is far from clear what the optimal responses to these changes in our climate system are. One reason for this is that we are facing severe uncertainty regarding the physical facts about the phenomenon of climate change. The Intergovernmental Panel on Climate Change (IPCC) was established to synthesize the latest scientific knowledge on climate change and to communicate it to policy makers. To this end, the IPCC developed as a key element of its toolbox an uncertainty framework.

In this paper, I assess the latest version of the uncertainty framework. First, I look at the meta-documents which characterize the uncertainty framework. I argue that the framework suffers from substantial conceptual issues. Secondly, I focus on the full report of Working Group I, which discusses the physical science basis of climate change, to explore how the uncertainty framework is put into practice. I show that the conceptual problems of the framework manifest themselves in concrete practical problems for the authors of the assessment report. Based on these observations, I suggest, thirdly, improvements to make the uncertainty framework more fruitful in the context of climate policy-making and, potentially, in other areas of policy-making as well.

The uncertainty framework equips the scientist with a confidence and a likelihood metric to qualify her statements about the causes and effects of climate change. An example for its application is the following statement:

“In the Northern Hemisphere, 1983-2012 was likely the warmest 30-year period of the last 1400 years (medium confidence)” (IPCC 2013, 3)

Confidence describes the validity of a finding. The confidence judgement is the result of the aggregation of the two sub-metrics ‘evidence’ and ‘agreement’. The evidence judgment is the result of an aggregation across different dimensions: type of evidence, its amount, its quality, and its consistency. The agreement judgment is capturing the consensus across the scientific community on a given finding. The likelihood metric is a quantified measure of uncertainty and expresses a probabilistic estimate of the occurrence of events or outcomes. The verdict ‘likely’, for example, is associated with a probability range of 66-100%.

The discussion of the uncertainty framework in the IPCC’s meta-documents reveals fundamental conceptual problems. The problems can be grouped into three categories: problems associated with (a) the lack of definitions of key terms, (b) the lack of specifications of relations between key terms, and (c) epistemological assumptions. Subsequently, I highlight a selection of problems. In relation to category (a), the agreement sub-metric is defined both as the degree of consensus between scientific publications on a finding as well as the number of competing causal explanations for a finding. Furthermore, the dimension ‘quality of evidence’ is not specified at all. In relation to category (b), the relationship between the sub-metric ‘agreement’ and the dimension ‘consistency of evidence’ is not clear because consistency is introduced as the degree to which evidence supports single or multiple explanations or projections. In addition, the meta-documents support two, mutually inconsistent interpretations of the relationship between the confidence and the likelihood metric. A first interpretation sees confidence statements as meta-judgment about the validity of the finding whereas likelihood statements are, uncorrelated, intra-finding judgments about the probability of an event or outcome described in the finding. A second interpretation understands confidence statements and likelihood statements as conveying the same information, whereas confidence judgments are used for qualitative and likelihood judgments are used for quantitative evidence. In relation to category (c), the fact that robust evidence (i.e. multiple lines of high quality, independent, and consistent evidence) could, according to the uncertainty framework, appear in combination with low, medium, or high agreement in the scientific community is puzzling from an epistemological point of view.

These conceptual problems give rise to concrete practical issues for the authors of the assessment report. For example, the authors fill the lacuna which is generated by the absence of the specification of the relationship between the agreement sub-metric and the dimension ‘consistency of evidence’ by equating these two terms. Moreover, the absence of a specification of the dimension ‘quality of evidence’ and the evidence categories ‘limited’ and ‘medium’ gives rise to a large ambiguity in the application of the uncertainty framework. The authors have to make decisions and there is no indication present in the report that these decisions are taken in a consistent and non-arbitrary way.

Based on these observations, I identify the construction of the confidence metric as the key issue which needs attention. I motivate three changes to the confidence metric. First, the agreement sub-metric should be interpreted as consistency of evidence indicating whether the set of evidence supports qualitatively agreeing or disagreeing estimates of parameters. Secondly, the evidence sub-metric should be reduced to an aggregation of the dimensions ‘quality’, ‘amount’, and ‘independence’ of evidence whereat quality of evidence has to be spelled out differently for different categories of evidence (e.g. mechanistic understanding, observational data, and model results). I use Douglas (2013)'s fine-grained typology of epistemic values, which distinguishes between minimum requirements and desiderata in this set of values, to give indications of defensible quality criteria. Thirdly, all categories of the two sub-metrics (e.g. robust evidence or medium agreement) should be introduced with paradigmatic examples. Based on this reworking of the two dimensions 'evidence' and 'agreement', I suggest an alternative aggregation mechanism into overall confidence judgements which places additional weight on the evidence metric. I close by arguing that the new confidence metric enables us to specify the relationship between the confidence and the likelihood metric in a straightforward way.

These three changes do not only provide the IPCC with a conceptually more coherent uncertainty framework but also allow exploring whether this confidence metric can be used in other policy contexts. One example is macroeconomics where we face complex systems with expectation-driven feedback mechanisms and considerable disagreement about the relevant causal structures.

Lena Zuchowski. A sideways glance at Smale's fourteenth problem: Definition and ontology of chaos

I will use Smale's fourteenth problem as a starting point for a discussion of the definition and ontology of chaos.

Thereby, I will first address the question whether the definition of chaos that has been applied to Smale's horseshoe map is one that also fits the Lorenz equations. Through analysing the well-known logistic equation in three different ways, I show that, in mathemat- ical practice, there exist two different definitions of chaos: periodic and aperiodic chaos. I then maintain that the chaos in Smale's horseshoe map is periodic, while that in the Lorenz equations is aperiodic.

Secondly, will survey the systems used to justify the definitions of periodic and ape- riodic chaos and categorized them as being continuous (C-Models) or discrete (D-Models). The results from this classification exercise indicate that chaos should be viewed as a prop- erty of a specific model rather than a property of a system of equations. I have also identified three trends in the way the types of models and chaos correlate: namely, that chaos ap- pears only in iterative modes of analysis; that aperiodic chaos is most often a property of discrete models; and that periodic chaos is related to continuous models. My analysis raises a number of worries that (I will argue) should receive more attention from philosophers of science.

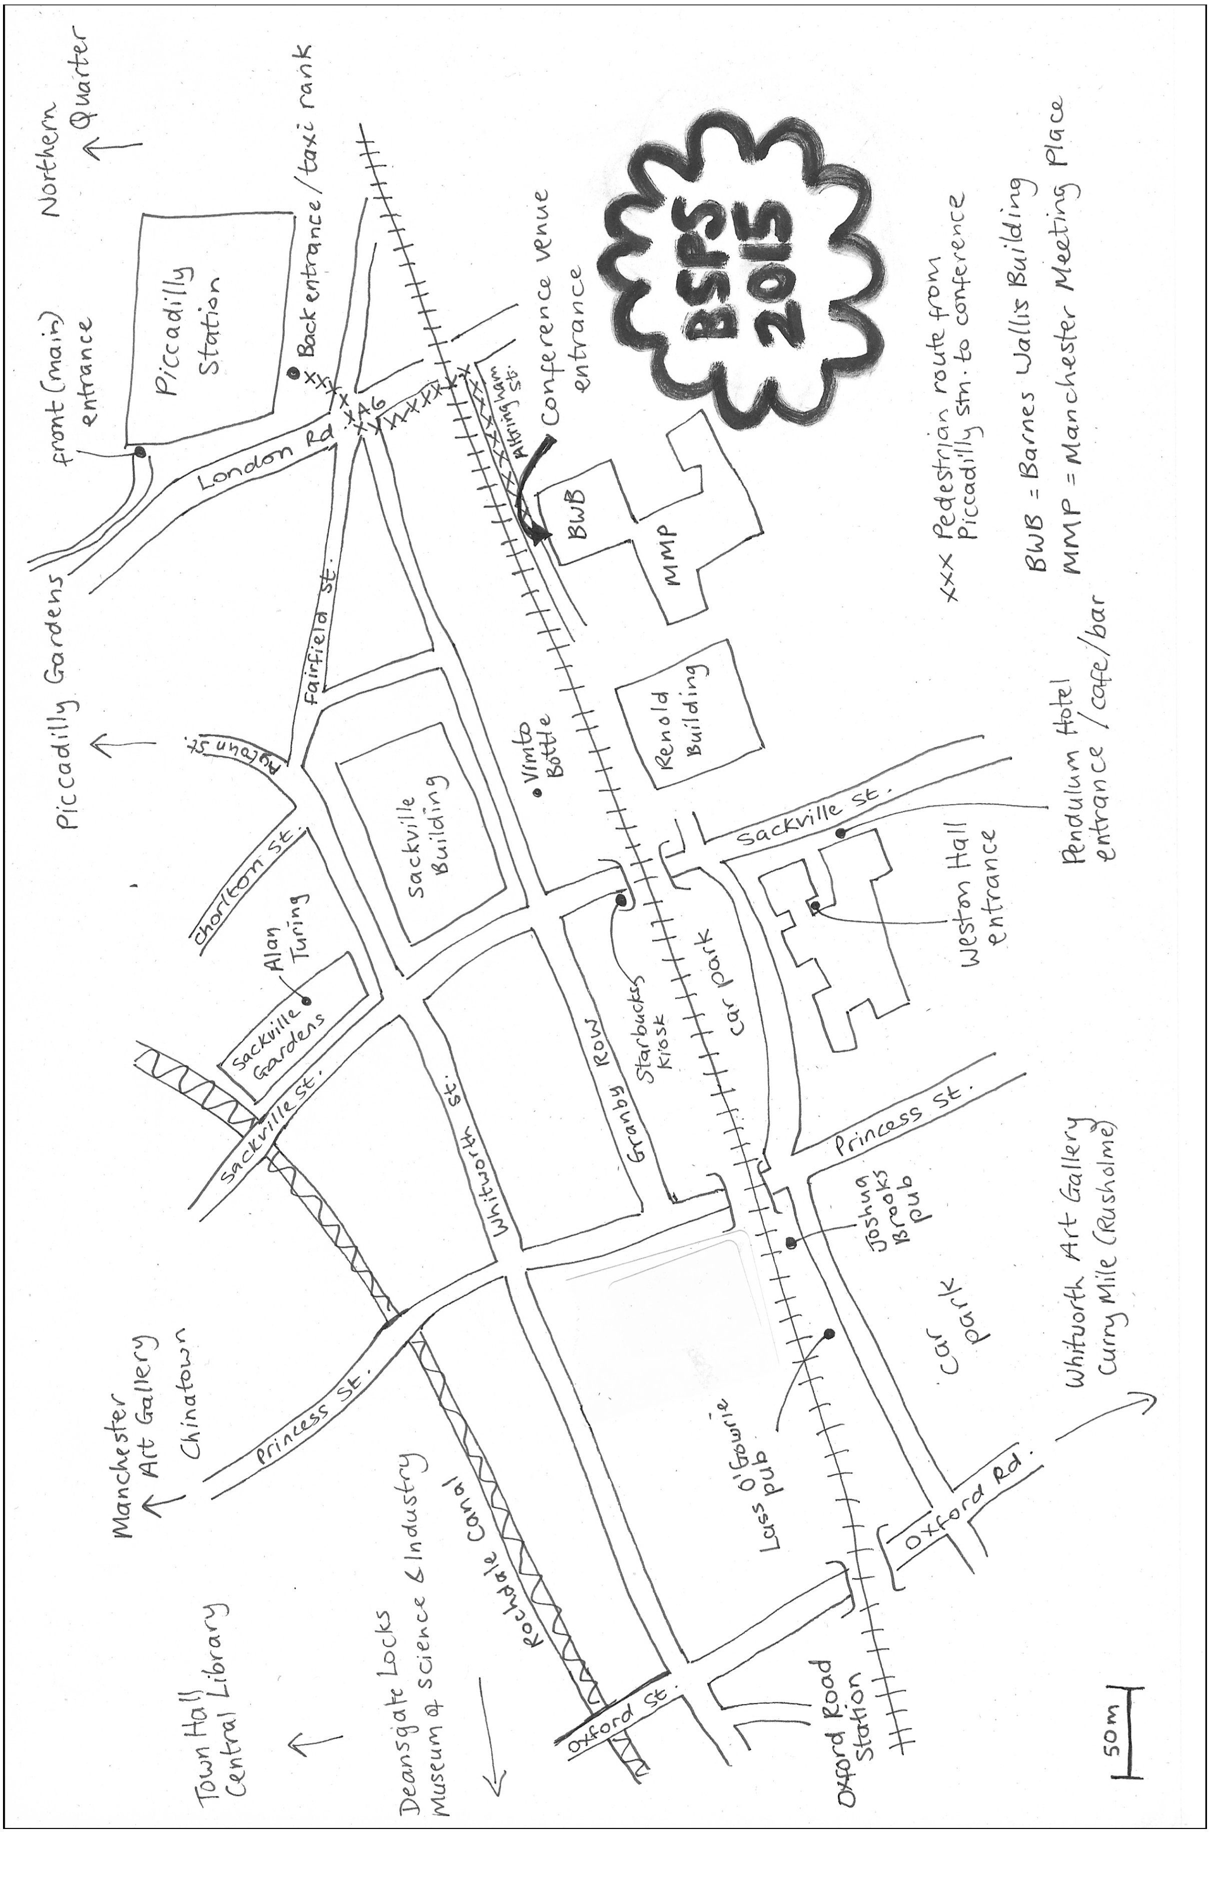

BSPS 2015
Northern Quarter
Piccadilly Station
front (main) entrance
Back entrance / taxi rank
London Rd.
A6
Altringham St.
Conference venue entrance
BWB
MMP
Renold Building
Fairfield St.
Piccadilly Gardens
Aytoun St.
Chorlton St.
Sackville Building
Vimto Bottle
Sackville St.
Pendulum Hotel entrance / cafe/bar
Weston Hall entrance
Alan Turing
Sackville Gardens
Sackville St.
Starbucks kiosk
car park
Granby Row
Whitworth St.
Princess St.
Joshua Brooks pub
car park
Whitworth Art Gallery
Curry Mile (Rusholme)
Lass O'Gowrie pub
Oxford Rd.
Oxford Road Station
Princess St.
Manchester Art Gallery
Chinatown
Rochdale Canal
Deansgate Locks
Museum of Science & Industry
Town Hall
Central Library
Oxford St.
xxx Pedestrian route from Piccadilly stn. to conference
BWB = Barnes Wallis Building
MMP = Manchester Meeting Place
50m