

**BSPS Annual Conference 2014**  
University of Cambridge

Fitzwilliam College, 10–11 July 2014



# Timetable

Map references appear in square-brackets (see back of programme).

Publishers' stands are in the Upper Hall 2 and Auditorium throughout the conference.

The Upper Hall used to be known as the 'Old Library'; both names may be in use on signs within the College.

## Thursday 10 July

09.00–10.00	Registration	Auditorium	[1A]
10.00–11.30	Plenary Session I PAUL GRIFFITHS (Sydney; Exeter) <i>Causation and information in living systems</i> Chaired by John Dupré (Exeter)	Auditorium	[1A]
11.30–12.00	Tea and Coffee	Upper Hall 2	[3A]
12.00–13.00	Open Session I	see p.3	
13.00–14.00	Lunch	Hall	[2D]
14.00–15.30	Open Session II	see p.3	
15.30–16.00	Tea and Coffee	Upper Hall 2	[3A]
16.00–17.30	Open Session III	see p.3	
17.45–19.15	The Presidential Address PETER CLARK (St Andrews) <i>Logic, applied mathematics and intuition</i> Chaired by Steven French (Leeds)	Auditorium	[1A]
19.15–20.00	Drinks Reception, sponsored by OUP	Grove Lawns	
20.15–21.30	Conference dinner	Hall	[2D]
20.00–00.00	Bar open	Café / Bar	[2E]

## Friday 11 July

07.30–08.30	Breakfast	Hall*	[2D]
08.45–10.15	Open Session IV	see p.4	
10.15–10.30	Tea and Coffee	Auditorium	[1A]
10.30–12.00	Plenary session II LAURA RUETSCHÉ (Michigan) <i>'Naturalistic' Metaphysics and the Interpretation of Quantum Theories</i> Chaired by Harvey Brown (Oxford)	Auditorium	[1A]
12.00–12.45	BSPS AGM	Auditorium	[1A]
12.45–13.45	Lunch	Hall	[2D]
13.45–14.45	Open Session V	see p.4	
15.00–16.30	Plenary Session III CHRISTOPHER PINCOCK (OHIO State) <i>Inference to the Best Explanation: A Modest Proposal</i> Chaired by Peter Clark (St Andrews)	Auditorium	[1A]

\*Accommodation includes breakfast at, and only at, the institution where accommodation is provided. So: those with accommodation at Murray Edwards must eat breakfast at Murray Edwards; those with accommodation at Fitzwilliam must eat breakfast at Fitzwilliam.

## Open sessions: Thursday

	Old SCR [2B]	Walter Grave Room [2C]	Gaskoin Room [3C]	Music Room [3D]	William Thatcher [1C]
<b>Session I</b> 12.00–12.30	SEBASTIAN LUTZ. Abstraction, idealization, and the application of mathematics	ADAM TOON. Where is the understanding?	MATTHEW PARKER. The poverty of infinitesimal probabilities	ELLEN FRIDLAND. Intelligence automatically	MATTHIAS EGG. Views of the quantum state in Bohmian Mechanics and the GRW Theory
12.30–13.00	ZEE PERRY. Intensive and extensive quantities	MARIO SANTOS-SOUSA. What, if anything, can the epistemology of number learn from the psychology of numerical cognition?	ALEXANDRU MARCOCI. Solving the absentminded driver problem through deliberation	KATHARINA KRAUS. Quantifying introspection? – The case of pain measurement	BENJAMIN EVA. Interpreting Topos Quantum Theory
<b>Session II</b> 14.00–14.30	PHILIPPE VERREAULT-JULIEN. Understanding through counterfactual analysis modelling	ALASTAIR WILSON. Towards a hybrid theory of laws	MAURICIO SUÁREZ. Probabilistic dispositions, chance distributions, and experimental statistics	GRANT RAMSEY & CHARLES PENCE. Is organismic fitness at the basis of evolutionary theory?	SAMUEL FLETCHER. Global spacetime similarity
14.30–15.00	CHRISTOPHER CLARKE. How economists' models of decision-making explain (even when false)	JOHN ROBERTS. Humean laws and explanation	NICK TOSH. Reviving finite frequentism: Humean chance without best systems	TIM LEWENS. The perils of cultural models	JULIUSZ DOBOSZEWSKI & TOMASZ PLACEK. Determinism and initial value problem in general relativity
15.00–15.30	ROBERT NORTHCOTT & ANNA ALEXANDROVA. Armchair Science	FOAD DIZADJI-BAHMANI & SEAMUS BRADLEY. Lewis' account of counterfactuals is incongruent with Lewis' account of laws of nature	HARJIT BHOGAL. Chance and Explanation: Why the New Principle is false	MAKMILLER PEDROSO. The evolution of transient individuals	J. BRIAN PITTS. Real change in Hamiltonian General Relativity
<b>Session III</b> 16.00–16.30	LAURA FELLINE. Causation, regularities and counterfactuals in fundamental physics: a solution to the bottoming-out problem	TOBY FRIEND. Laws as analysans for causation	TEDDY GROVES. Accuracy arguments in the context of Carnapian inductive logic	ELSELIJN KINGMA. Metaphysics of pregnancy: Fetuses as part of the maternal organism	RUNE NYRUP. Analogical reasoning and pursuitworthiness
16.30–17.00	LORENZO CASINI & JON WILLIAMSON. How to model mechanisms	ANDREAS HÜTTEMANN. Actual causation and default processes	JÜRGEN LANDES. Strictly proper scoring rules and the Probability Norm	ARGYRIS ARNELLOS. An organizational account of organismically integrated wholes	RADIN DARDASHTI, KARIM THEBAULT & ERIC WINSBERG. Confirmation via analogue simulation: What dumb holes can tell us about gravity
17.00–17.30	MATTEO COLOMBO & STEPHAN HARTMANN. Bayesian cognitive science, unification, and explanation	PETER FAZEKAS, BALÁZS GYENIS, GÁBOR HOFER-SZABÓ & GERGELY KERTÉSZ. A dynamical systems approach to causation	DANIEL MALINSKY. Hypothesis testing, 'Dutch Book' arguments, and risk	TERO IJÁS. Beyond tinkering: Design and understanding through directed evolution in synthetic biology – a case from protein design	ERIK CURIEL. Carnot Cycles and black hole entropy

## Open sessions: Friday

	Old SCR [2B]	Walter Grave Room [2C]	Gaskoin Room [3C]	Music Room [3D]	William Thatcher [1C]
<b>Session IV</b> 08.45–09.15	DAVID TEPLow. Alzheimer's disease: Philosophical impediments towards a cure	ALEXANDER REUTLINGER. What's explanatory about non-causal explanations?	ALESSANDRA BASSO. The triangulation of measurement procedures	CHARLOTTE WERNDL. On defining climate and climate change	CONOR MAYO-WILSON. Structural chaos
09.15–09.45	DANA TULODZIECKI. The pessimistic meta-induction and the superfluity of approximate truth	JUHA SAATSI. Worthwhile distinctions: Kinematic, dynamic, and (non-)causal explanations	JAAKKO KUORIKOSKI & CATERINA MARCHIONNI. Evidential diversity and the triangulation of phenomena	CARLO MARTINI. The limits of trust in interdisciplinary science	LENA ZUCHOWSKI. Revisiting Smale's 14th problem: Are there two kinds of chaos?
09.45–10.15	BON-HYUK KOO. How much can we grasp? Objective blind realism as an answer to pessimistic meta-induction and Stanford's 'trust' argument	M. CHIRIMUUTA. Efficient coding explanations in neuroscience: Causal and non-causal	CHIARA LISCIANDRA. Robustness analysis as a non-empirical confirmatory practice	PIOTR SZALEK. The Duhem–Quine Thesis reconsidered	JAMES FRASER. Spontaneous symmetry breaking in finite systems
<b>Session V</b> 13.45–14.15	DAVE RACE. Filling in surplus structure in the partial structures framework	JONATHAN BAIN. What explains the spin–statistics connection?	LEE ELKIN. A conciliation model for polarized beliefs	ARIANNE SHAHVISI. Eliminating conspiracies via the genealogy of subsystems	CRAIG CALLENDER & CHRISTIAN WUTHRICH. What becomes of a causal set?
14.15–14.45	JAMES NGUYEN. Why data models do not supply the target structure required by the structuralist account of scientific representation	RYAN SAMAROO. There is no conspiracy of inertia	DAVID GLASS & MARK MCCARTNEY. Explanatory competition and explaining away	ADAM WHITE. Emergence in biological pathways	CARLO ROSSI. enduring a relativistic world

**Argyris Arnellos. An organizational account of organismically integrated wholes** From an organizational perspective, organisms should not only be capable of reproducing each of their own differentiated parts but also the dynamic and functional interrelationships between those parts, i.e. their own global/collective organization. Moreover, apart from their constructive dimension, organisms are also agents engaging in interactions with their environments, in a way that these interactions are in a functional and reciprocal relation (at least) with the maintenance of their global organization. Then, one should not focus on how aggregations of parts become temporary cohesive systems, but on how they may turn into the respective highly organized and functionally integrated and differentiated wholes that adaptively interact with their environments.

This is quite challenging, especially with respect to biological organisms, where the concept of functional integration is often accused of looseness that allows for an excessive plurality of collaboratively-produced heterogeneous organismal wholes. Indeed, from the early stages of collaboration in the biological world, entities assemble into groups, bringing forth several types of relatively stable cellular associations (e.g. biofilms, filaments, colonies, various types of aggregations, pluricellular systems, modular systems, etc). All these aggregation comprise a number of different cell types (though relatively low) and they are characterized by specialized intercellular interactions, thereby exhibiting a degree of functional integration. In turn, the result of this integration (the various functional interactions between the cells) is externally observed, at least, as a global agential activity through which the group expands its overall adaptive capacity.

I'll begin by explaining why the minimization or even the complete elimination of the possibility of conflicts between the parts (alignment of fitness) together with the achievement of a clear and functional division of labor (export of fitness) are not enough for organismal wholes, since, notwithstanding the underlying integration in such cases the agential dimension is not satisfied. I will then suggest a general scheme of organizational conditions and requirements for the realization of the special kind of functionally integrated differentiation necessary for organisms. More specifically, I will argue that an organismal whole is the result of an endogenously produced regulatory logic, whose various operational patterns control both the generation and the integration of the functionally differentiated parts as well as the interactive behavior of those constitutive parts, so that the ensemble becomes a functionally cohesive self-maintaining/reproducing organization capable of adaptive interaction with its environment. I will explain the structural and operational characteristics of this regulatory logic showing that organismal wholes are not just the result of generation of functional diversity, but that what is more important is its control through the regulatory relationships among the (increasingly complex) components and modules of the system.

**Jonathan Bain. What explains the spin-statistics connection?** The spin-statistics connection (SSC) plays an essential role in explanations of non-relativistic quantum phenomena such as the electronic structure of solids and the behavior of Bose condensates and superconductors. However, it is only derivable in the context of relativistic quantum field theories (RQFTs) in the form of the Spin-Statistics theorem; and there are mutually incompatible ways of deriving it. This essay considers the sense in which SSC is an essential property in RQFTs, and how it is that an essential property in one type of

theory can appear in fundamental explanations offered by other, inherently distinct theories.

The first part of the essay argues that an explanation of SSC based on the Spin-Statistics theorem is best understood as structural in the following sense: the Spin-Statistics theorem demonstrates how a set of principles, the contents of which is specific to a particular approach to RQFTs, limits the admissible states of physical systems described by that approach to those that possess SSC. This way of understanding the spin-statistics connection cannot be formulated in terms of DN, unifying, or causal/mechanical explanations.

The second part of the essay argues that a structural explanation of SSC is problematic because (a) there are different ways of formulating the Spin-Statistics theorem that disagree on the principles essential to the derivation; and (b) SSC plays a fundamental role in many explanations in non-relativistic quantum mechanics (NQM) and non-relativistic quantum field theories (NQFTs). This is puzzling to the extent that relativity is essential to explanations of SSC in RQFTs: Why is an (apparently) essentially relativistic property fundamental to explanations of some non-relativistic physical systems?

Thus a full account of SSC should explain by virtue of both a derivation in RQFTs, and an appeal to intertheoretic relations between RQFTs on the one hand, and NQFTs and NQM on the other. The third part of the essay compares this type of explanation with a similar account given by Weatherall (2011). Weatherall offers an explanation of a feature of the world (the equality of gravitational and inertial mass) that is expressed in one theory, Newtonian gravity, and that can only be adequately understood by appealing to another, presumably more fundamental theory, general relativity (GR). The explanatory work is done by means of a translation between GR and Newtonian gravity. In the present essay, an explanation is given of a feature of the world (SSC) that is expressed in one type of theory, NQM and NQFT, and that can only be adequately understood by appealing to another, presumably more fundamental theory, RQFT. As in Weatherall's example, the explanatory work is done (in part) by means of a translation between theories. However, the role of the translation in both cases differs: In the spin-statistics case, the translation is essential to the explanation insofar as it demonstrates how the explanandum is a consequence (in part) of an essential property (Lorentz invariance) of the fundamental theory, and yet also appears essentially in the less fundamental theories that are not characterized by this property.

**Alessandra Basso. The triangulation of measurement procedures.** Philosophers of science have discussed measurement triangulation as an exemplary case in which the appeals for triangulation are normatively good arguments. Different and independent measurement procedures, none of which is clearly superior to the others, can be used to triangulate on the targeted property. The agreement of different measurements increases our confidence in the results by guarding against errors in the procedure. These arguments are intuitively appealing, but several questions remain unexplored about the use of triangulation in this context.

Observation of measurement assessment practice reveals that scientists do not employ substantially different measurement procedures for testing their measurements, but rather focus on repeating measurement under similar conditions or under controlled variations of the procedure. In other words, scientists seem to assume that entirely different procedures necessarily lead to incompatible measurement results.

Recent philosophical works on measurement accommodate this observation. Hasok Chang, for instance, emphasizes that, prior to the construction of an accepted measurement procedure, there is no (good) evidence of the target and hence no other (good) epistemic access to it. Bas van Fraassen, moreover, argues that measurement does not show what the target is like ‘in itself’ but how it looks like under the specific measurement design, and therefore substantially different procedures generally lead to incompatible results. Furthermore, it has been claimed that it is impossible to develop distinct and entirely independent procedures for measuring the same target. Any measurement procedure must be based on the current theoretical knowledge about the targeted property and its interaction with the surrounding environment and hence it is impossible to have entirely independent ways of determining the same property, because they always share a common background theory.

In reply to these worries, it is possible to maintain that complete independence is not required for triangulation, and advocate a weaker definition of independence – as it has recently been suggested in the context of the robustness analysis of models. This line of argument is promising, but in order to make it more precise, it is necessary to specify what kind of independence is required for the assessment of measurements.

The criteria for independence, however, appear to be substantially different across disciplines. Different disciplines use distinct accounts of measurement quality, different ways of testing them and different thresholds for what is considered sufficient independence.

The discipline-specific accounts of independence can deal with the different practical problems that scientists face in their own disciplines, and hence the fragmentation of this literature raises interesting questions for interdisciplinary studies. What are the reasons of these differences? This paper addresses this question by investigating and comparing the practice of measurement assessment in different disciplines. The investigation is based on discipline-specific guidelines for the assessment of measurement and on the observation of measurement assessment practice. I argue that, perhaps surprisingly, the social sciences tend to have more demanding criteria than the natural sciences.

**Harjit Bhogal. Chance and explanation: Why the New Principle is false.** Hall (2004) argues that the way chance constrains rational credence is given by the New Principle. Roughly, the idea is that we should set our credence equal to our expectation of the chance of A conditional on our evidence.

This view is naturally motivated by a view of chance where chance is an analyst expert. An analyst expert is someone who correctly evaluates the force of evidence. To use Rachel Briggs’ (2009) example, a good advice columnist is an analyst expert. They do not have more evidence than you, but they are better at judging what is correct given the evidence.

In this paper I argue that the NP is false by considering a case where chance fails to be analyst expert – that is, a case where it fails to correctly evaluate the evidence it has.

The central idea is that chances explain events. And they explain by encoding explanatory information. This leads to a problem when we have information that is evidentially, but not explanatorily, relevant to the occurrence of an event. Chance must ‘ignore’ such information otherwise it would fail to be explanatory. But in ignoring this information it fails as an analyst expert. I give a crystal ball case which illustrates this. In the case I give information about the output and reliability of the crystal ball is

evidentially but not explanatorily relevant to the event in question. It appears NP gets the wrong result in this case. I then consider, and reject, responses designed to show that NP is not to blame. Along the way I defend a principle about when chance is defined roughly that the chance of A relative to X only exists if X would explain A if both X and A held. This principle is used to cash out the idea that chance must ‘ignore’ non-explanatory information.

I then suggest an admissibility clause, formulated in terms of explanation, that can be added to the NP to get the right result.

Finally, I suggest that the considerations that have some before apply very neatly to the case of the interaction of special science chances. From the perspective of a certain special science chance fundamental chances are analogous to crystal ball information. Also special science chances ‘ignore’ non-explanatory information. And we can deal with the interaction of such chances with credence by using the NP with the admissibility clause I suggested.

### **Craig Callender & Christian Wuthrich. What becomes of a causal set?**

Contemporary physics is notoriously hostile to an A-theoretic metaphysics of time. A recent approach to quantum gravity promises to reverse that verdict: advocates of causal set theory (CST) have argued that their framework is at least consistent with a fundamental notion of ‘becoming’. How can a fundamental physical theory which claims to be fully relativistic and which aspires to describe structures that give rise to relativistic spacetimes support substantive becoming? The well-rehearsed difficulty here is that a global sense of objective becoming that has some hope of underwriting the usual A-theoretic motivations seems flatly incompatible with the relativity of simultaneity (and consequently the Lorentz symmetry) upheld in contemporary physics. We take this dilemma to be underwritten by a result in special relativity due to Howard Stein.

The analogue of Stein’s theorem does not hold in the context of CST. This fact gives renewed hope to the A-theorist to evade the dilemma, as it suggests that the A-theorist has more tools at her disposal in the context of CST than were available to her in relativistic physics. Unfortunately, this hope is not long lived and the A-theorist quickly finds herself facing the old dilemma again: the permitted relations are of a kind that cannot give rise to a robust notion of a macroscopic present or a form of becoming, on pain of violating the Lorentz symmetry assumed to be valid at those scales. The remaining kinds of becoming are restricted to a purely local feature of objective physical reality. The resulting combination of localized becoming with a block universe is reminiscent of ‘worldline’ or ‘past-lightcone becoming’ in Minkowski spacetime.

However, there is a novel and exotic notion of becoming compatible with CST, and hence arguably with relativity. To find anything smacking of becoming, one needs to turn to the theory’s dynamics. The dynamics for a causal set is a stochastic law of sequential growth. What grows are the number of elements. The Lorentz symmetry required in relativity gets transposed into a requirement that the sequential birthing occurs in an order that lacks any physical meaning apart from the fact that any event that causally precedes another must have been birthed earlier. While this rules out any determinate fact regarding which of two ‘spacelike’-related events was birthed first, it permits two objective and global ways in which there is temporal becoming. First, as a causal set grows from having N elements to have N + 1, although it may not be determinate which events have be-



come yet, it is a determinate fact that the number of events has increased by one. Second, although events may thus linger in this ontological penumbra for many stages of the temporal becoming, there will be a monotonically increasing number of events which have determinately become.

If it is coherent, therefore, to speak of a causal set having a certain number of elements but without saying what those elements are, then causal set theory does permit a new kind of—admittedly radical and bizarre—temporal becoming.

**Lorenzo Casini & Jon Williamson. How to model mechanisms.** Mechanisms are usually viewed as inherently hierarchical, with lower levels of a mechanism ‘constituting’ its higher-level behaviour, and the higher-level behaviour being ‘decomposable’ into lower-level entities and activities. The distinction between different levels is common to biological sciences, where macro-level features of the system (e.g., traits and functions), are explained in terms of properties and relations of parts (e.g., genes and proteins). In biology textbooks, verbal and pictorial descriptions of mechanisms are typically qualitative. It is often desirable to associate to such qualitative descriptions also a quantitative description, in order to facilitate explanatory and predictive tasks involving the complex relations across the levels. However, most available quantitative descriptions of biological mechanisms (e.g., differential equations, Petri nets, neural networks, Bayesian networks) fail to capture the hierarchical aspect of mechanisms.

To remedy this deficiency, the Recursive Bayesian Network (RBN) formalism was put forward by Casini et al. (2011) and its applicability extended to cyclic mechanisms by Clarke et al. (2013). In a nutshell, an RBN represents hierarchical relations by decomposing certain higher-level variables into lower-level causal graphs. The associated probability distribution must satisfy not only the causal Markov condition, as in traditional causal BNs, but also an additional condition, viz. the recursive Markov condition. Interlevel causal inferences are drawn with the aid of the probability distributions in the so-called ‘flattening’.

In this paper, we illustrate a further advantage of RBNs, namely how RBNs can be used to represent non-modular mechanisms involving cycles, which are resistant to modelling by means of DAGs. RBNs represent non-modularity in terms of decompositions of higher-level variables into overlapping lower-level causal graphs. Such complex relations are common to biological mechanisms, where lower-level entities are often involved in more than one higher-level function for the same behaviour. We illustrate this procedure with reference to the internal pathway for apoptosis, where two overlapping feedbacks cooperate to making apoptosis irreversible (Legewie et al., 2006). In particular, we show how interlevel inferences are drawn between higher-level variables on the one hand, and lower-level variables decomposing non-modular functions on the other. To this end, one needs knowledge of the relevant conditional probabilities in the flattening. If not directly inferrable from available datasets, these are calculated by selecting the probability distribution that, among those that satisfy the RBN constraints (conditional independences and conditional probabilities), maximises entropy (Williamson, 2010).

Finally, we suggest that the applicability of the notion of mechanistic decomposition depends (among other things) on the degree of modularity: the larger the constitutional overlap, the less the distinction between entities and functions at different levels makes sense.

Casini, L., Illari, P. M., Russo, F., and Williamson, J. (2011). Models for Prediction, Explanation, and Control: Recursive Bayesian Networks, *Theoria*, 70:5-33

Clarke, B., Leuridan, B., and Williamson, J. (2013). Modeling Mechanisms with Causal Cycles, *Synthese*, doi:10.1007/s11229-013-0360-7

Legewie, S., Blüthgen, N., and Herzog, H. (2006). Mathematical Modeling Identifies Inhibitors of Apoptosis as Mediators of Positive Feedback and Bistability, *PLoS Computational Biology*, 2(9):1061–1073.

Williamson, J. (2010). In Defence of Objective Bayesianism, Oxford University Press.

**M. Chirimuuta. Efficient coding explanations in neuroscience: Causal and non-causal.** In a recent paper (Author) I argue that efficient coding explanations in computational neuroscience are distinct from mechanistic explanations, and that they have an important role to play in the development of theories of neural coding. Efficient coding explanations account for the observed properties of neural circuits in terms of the computational advantages of particular arrangements of neurons, appealing to coding principles such as redundancy reduction (Attneave 1954, Barlow 1961) and decorrelation (Schwartz and Simoncelli 2001).

For example, it has long been observed that neurons in primary visual cortex (V1) have elongated receptive fields (RF’s), and are particularly responsive to bar-like stimuli of a particular width and orientation (Hubel and Wiesel 1962). These RF’s are commonly modelled by a two dimensional Gabor function (a sinusoidal function combined with a Gaussian envelope). An important question for theoretical visual neuroscience is, why do V1 neurons have receptive fields that can be fit by the Gabor equation? In the paper I discuss two different explanations which have been proposed in response to this question, arguing that one is a form of non-causal explanation and that the other is kind of (non-mechanistic) causal explanation.

Firstly, it has been observed that the Gabor equation has interesting information-theoretic property of minimising joint uncertainty over spatial position and width of bar-like stimuli (Daugman 1985). In other words, given that there is an inherent trade-off between knowing the spatial position and the width of the stimuli, the Gabor function is a theoretically optimal filter for recovering both types of information. Such trade-offs are brute mathematical facts about the universe, and so cannot be intervened on (even hypothetically). However, we do explain features of the world in terms of these mathematical facts, because, in Woodward’s terms, they address what-if-things-had-been-different questions. In Woodward’s (2003) example of structural/mathematical explanation, the stability of planetary orbits depends counterfactually on the 4D structure of space time. I argue that this example is analogous to Daugman’s (1985) explanation of the RF properties. I also discuss this example in the light of Rice’s (in press) account of non-causal optimality explanation in biology and economics.

Secondly, Hyvärinen and Hoyer (2001) have proposed that, ‘[t]he reason why the RFs have Gabor-like shapes might thus be that this kind of RFs are optimal for analyzing the input that the visual system typically receives.’ I argue that this is a kind of causal explanation which is analogous to Mayr’s (1961) notion of ultimate causal explanation and that it also fits into Potochnik’s (2010) and Elgin and Sober’s (2002) censored-causal framework for understanding optimality explanation in biology.

Importantly, such explanations support intervention interpretations because we can alter the developmental environment of organisms and look for resulting changes in RF structure (Blake-moore and Cooper 1970, Wainwright et al 2001).

To conclude, I discuss how both the causal and non-causal kinds of explanation have played a role in guiding experimental research within visual neuroscience.

**Christopher Clarke. How economists' models of decision-making explain (even when false).** Most economists think that the content of economic models of decision-making is not what it appears to be at face value (Friedman and Savage 1952; Binmore 2009). Economists take such models to describe more or less accurately the choices that economic agents make; but not to describe the cognitive process through which agents make their choices; nor to describe an agent's commonsense beliefs and desires. Many philosophers of economics disagree with this instrumentalist story about the content of economic models (Craver and Alexandrova 2008; Hausman 2012).

I will not evaluate this instrumentalist story. Instead I will ask: what are the consequences of this instrumentalist story, if true? The consensus is that instrumentalist models of decision-making do not explain (Bermudez 2009). On this point one sees agreement between those who oppose the instrumentalist story (Robbins 1932, Rosenberg 1992) and those who endorse it (Samuelson 1938, Binmore 2005).

This paper rejects this consensus position: instrumentalist economic models of decision-making sometimes do provide deep explanations. My example will be microeconomic models of market equilibrium. Importantly, I will make this case whilst maintaining a high standard for what counts as an explanation: to explain is to provide knowledge of the causes of a phenomenon. (Contrast this high standard with recent attempts to rescue economic models as merely 'how possibly' explanations.)

I finish by squaring my view with Pettit's (1995) view on economic explanations and with recent discussions of whether game theoretic models explain (Alexandrova and Northcott 2014).

**Matteo Colombo & Stephan Hartmann. Bayesian cognitive science, unification, and explanation.** A recurrent claim made in the growing literature in Bayesian cognitive science is that one of the greatest values of studying phenomena such as perception, action, categorization, reasoning, learning, and decision-making within the Bayesian framework consists in the unifying power of this modelling framework.

An assumption often implicit in this literature is that unification obviously bears on explanation. However, the link between unification and explanation is far from obvious, as unification in science is a heterogeneous notion, which may have little to do with explanation. So, it is not clear in which sense the kind of unification produced by Bayesian modelling in cognitive science is explanatory. If the relationship between unification, explanation, and Bayesian modelling in cognitive science is elucidated, then the debate over the virtues and pitfalls of the Bayesian approach will make a step forward. The goal of the present paper is to elucidate such a relationship.

After an overview of the Bayesian framework and of the variety of phenomena recently studied within this framework in cognitive science, we ask: How is unification produced within Bayesian cognitive science? To address this question, we focus

on the case of cue combination. This case illustrates how diverse phenomena can be unified within the Bayesian framework. It will help us to argue that unification in Bayesian cognitive science is driven by the mathematics of Bayesian decision theory, rather than by some causal hypothesis concerning how different phenomena are brought about by a single type of mechanism.

As there is no agreement on cases or accounts of genuine explanation, we shall not assume that Bayesian unification necessarily contributes (or fails to contribute) explanatory power. We shall focus our attention on the relationship between Bayesian unification and causal-mechanical explanation, assuming that one prominent feature of many adequate explanations of cognitive phenomena is that they reveal at least some relevant aspects of the mechanisms that produce those phenomena. Given this plausible assumption, the second question we ask is: What types of constraints can Bayesian unification place on causal-mechanical explanation in cognitive science? We shall address this question, showing that some features of Bayesian unification can play at least a heuristic role in the discovery and confirmation of the mechanisms of some cognitive phenomena. If these heuristics contribute to revealing some relevant aspects of the mechanisms that produce phenomena of interest, then Bayesian unification has genuine explanatory traction.

Our novel contributions to existing literature are twofold. First, Bayesian unification is not obviously linked to causal-mechanical explanation: unification in Bayesian cognitive science is driven by the mathematics of Bayesian decision theory, rather than by some causal hypothesis concerning how different phenomena are brought about by a single type of mechanism. Second, Bayesian unification can place fruitful constraints on causal-mechanical explanation. Specifically, it can place constraints on mechanism discovery, on the identification of relevant mechanistic features, and on confirmation of competitive mechanistic models.

**Erik Curiel. Carnot Cycles and black hole entropy.**

It is universally accepted in the physics literature that the striking formal analogy between the four laws obeyed by classical black holes and the four Laws of thermodynamics is just that—a formal analogy. In order to take the analogy seriously, and conclude that black holes are true thermodynamical objects, and that the laws they obey really are the laws of thermodynamics extended to treat them, one must take quantum effects into account, such as Hawking radiation. I argue that the standard arguments given in defense of this claim are not physically sound, and, in any event, beg the question.

In particular, the claim that classical black holes are "perfect absorbers", and thus have temperature absolute zero, is incorrect: when radiation or ordinary matter passes through the event horizon of a black hole, it does indeed emit 'energy', in the form of gravitational radiation. This gravitational radiation is the appropriate medium, moreover, for a characterization of "thermal coupling" between ordinary thermal systems and black holes. The standard argument, furthermore, begs the question in so far as it uses the definition of temperature derived from the theory of black-body radiation as the measure for the temperature of a classical black hole. That definition of temperature, however, is fundamentally a quantum one, in the sense that it requires a quantum theory—Planck's theory of black-body radiation—in order to formulate. A quantum theory of temperature, however, is not the appropriate one to use when one is considering whether or not to treat purely classical black holes as classical thermodynamical systems.

Thus I claim, to the contrary of the standard argument, that we should take the analogy between the laws of black holes and the laws of thermodynamics very seriously even when we restrict attention to the purely classical regime, and thus that we should think of the area of a classical black hole as a true thermodynamical entropy and the surface gravity as a true thermodynamical temperature in their own right, independent of their relation to quantum phenomena.

The strongest way to argue for this is to show that a black hole's area 'couples' in the appropriate way to the entropy, temperature and heat content of ordinary thermodynamical systems. Based on a mechanism proposed a long time ago by Robert Ge- roch, I accomplish this by constructing the appropriate analogue of a Carnot cycle that uses a classical Kerr black hole as one of the heat sinks, in which the entropy attributed to the black hole is exactly its Bekenstein-Hawking entropy, one-fourth its area (in natural units). I show that this allows one to define an absolute temperature scale for the black hole that quantitatively matches the one defined for the ordinary thermodynamical systems coupling to the black holes in the constructed Carnot cycle. According to the constructed scale, moreover, the absolute temperature of the black hole is exactly its Hawking temperature (in natural units), viz., its surface gravity multiplied by  $\frac{1}{4}\pi$ . In so far, therefore, as the black hole's area plays the physical role of an entropy in interactions with classical thermodynamical systems, and manifestly has all the properties of ordinary thermodynamical entropy, it is a fortiori a true physical entropy, for one can demand nothing else of a quantity in order to think of it so.

If this argument is correct, it shows that there are already deep connections between gravitation and thermodynamics in their own right, independent of any relation to or input from quantum field theory. Besides giving us insight into the character of classical general relativity and classical thermodynamics as theories, and the physical quantities they attribute to systems in the world, this would have profound consequences for our conceptual understanding of the broad class of phenomena both theories treat.

I conclude by discussing possible ways to extend the construction to more generalized purely gravitational systems for which it has been hypothesized and argued one should be able to attribute an entropy, such as causal horizons, gravitational radiation and singularities.

**Radin Dardashti, Karim Thebault & Eric Winsberg. Confirmation via analogue simulation: What dumb holes can tell us about gravity.** Philosophical analysis of science is always, at least partially, held hostage to the fortunes of scientific practice. As the ways in which scientists do science evolve, so must the models of science put forward by philosophers, else the discipline will inevitably decline into irrelevance. Here we will articulate a refinement and extension of existent analysis of the role of analogies in science inspired by fluid dynamical 'dumb hole' analogues to gravitational black holes. Our central claim, is that this case exemplifies a notion of analogue simulation that, unlike other species of analogical reasoning, can provide a conduit for confirmation. Trading on an exact syntactic isomorphism, analogue simulation allows certain inaccessible phenomenology in the target system to be probed by experimentation on the analogue. Given further model external and empirically grounded arguments, this then allows us to confirm the existence of novel phenomenology in the target system via the observation of its correlate in the analogue. The potential importance of this claim is particularly startling in the context of

our chosen example since Hawking radiation is among the gravitational phenomena that 'dumb holes' have the capacity to simulate, and by our lights confirm. Thus, if our analysis is correct, the quantum phenomenology of black holes is already within reach of contemporary experimental research in analogue gravity.

We will first briefly review the physical background necessary for a basic understanding of: Hawking radiation in semiclassical gravity; the modelling of sound in fluids; and the acoustic analogue model of Hawking radiation. We then give an explication of the idea of analogue simulation and our claim that it can provide a means for confirmation. First, we review the traditional notion of analogical reasoning, introduce a framework for understanding analogue simulation, and then contrast the two. These ideas are reinforced by consideration of a simple example of analogue simulation, based upon the connection between Coulomb and Newtonian gravitational forces, and told by means of a fable. This discussion leads naturally into the problem of justifying the inferences necessary for analogue simulation to enable confirmation. Our key idea is that in certain circumstances predictions concerning inaccessible phenomena can be confirmed via an analogue simulation in a different system. As we shall see, one is only justified in making such claims once one has established additional empirically grounded and model external reasons for the accuracy and robustness of the relevant modelling frameworks and syntactic isomorphism within the domains involved. The problems of experimental realisation of Hawking radiation and of finding such external reasons in the dumb hole/black hole case then become our main occupation. Before, in completion of our argument, we present the case for the dumb hole/black hole correspondence offering us the possibility for confirmation of Hawking radiation via analogue simulation. We conclude by offering a prospectus for extension of the idea of analogue simulation to other areas of science, and give a short sketch of one case of particularly obvious relevance.

**Foad Dizadji-Bahmani & Seamus Bradley. Lewis' account of counterfactuals is incongruent with Lewis' account of laws of nature.** In this paper we argue that there is a problem with the conjunction of Lewis' account of counterfactual conditionals and his account of laws of nature. This is a pressing problem since both accounts are individually plausible, and popular.

Lewis' account of counterfactuals involves appeal to similarity between possible worlds. Suppose we are interested in this counterfactual: 'If you had let go of the ball, it would have dropped to the floor'. Now consider the following possible worlds:

- (a) The actual world: you do not let go of the ball but hold it stationary. This is a  $\neg L \wedge \neg F$ -world.
- (b) A standard world where you let go of the ball and it falls to the ground. This is a  $L \wedge F$ -world.
- (c) A deviant world where you let go of the ball but it stays in the same position as it does in (a). This is an  $L \wedge \neg F$ -world.

At (c),  $L$  is true; (c) is a  $L$ -world. But we construct (c) in such a way as to be like the actual world – (a) – in all other respects. As we'll say, (c) is maximally similar excepting  $L$  to the actual world.

Under Lewis' account, for the counterfactual to be true requires (b) to be closer to the actual world, (a), than any world where I let go and the ball does not drop (any  $L \wedge \neg F$ -world).

However, (c) is a  $L \wedge \neg F$ -world which is closer to (a) than (b). Indeed (c) is maximally similar excepting  $L$  to the actual world. So under Lewis' account, the counterfactual comes out false, contrary to the intuition to which his account must do justice.

In response to this, Lewis offers an alternative account of similarity that appeals to laws of nature: closeness of worlds should be understood in terms of violations of the laws of nature, and in such terms the standard world is the closer to the actual world, he claims.

The problem is this: according to Lewis' account, laws of nature just are some sort of summaries of particular matters of local fact. But if that is right then the actual laws of nature and the laws of nature at the deviant world will be less different than the actual laws of nature and the laws of nature at the standard world, or so we argue.

To summarise, given Lewis' account of laws of nature underpinning the similarity measure between possible worlds, as per the hierarchy aforementioned, deviant worlds come out as closer to the actual world than standard worlds. Thus, standard worlds fail to ground the intuitively correct truth-values of counterfactuals. So, given the conjunction of Lewis' account of counterfactuals and his account of laws of nature, one ends up with the intuitively wrong truth-value assignments to counterfactuals. The central claim of our paper is, therefore, that his two accounts are incongruent.

### **Juliusz Doboszewski & Tomasz Placek. Determinism and initial value problem in general relativity.**

Physicists' concern with the determinism of general relativity (GR) comes from their interest in whether or not GR admits a globally well-posed initial value problem. This problem requires one to consider a 'slice' of a GR spacetime at a given time, with a data set on that slice, and ask if the data set uniquely determines a global solution.

In the talk we limit our attention to globally hyperbolic spacetimes (of dimension  $n$ ) that are vacuum solutions of the Einstein equations, that is, for which the Ricci curvature tensor vanishes. In this case a vacuum data set is a triple  $\langle D, g, k \rangle$ , (where  $D$  is an  $(n - 1)$ -dimensional manifold,  $g$  is a Riemannian metric, and  $k$  a symmetric covariant tensor), which satisfies certain equations called initial value constraints (see R. Wald's *General Relativity*, 1984).

The Choquet-Bruhat and Geroch theorem (Comm Math Phys 14, 1969) is relevant to whether or not GR is deterministic in the vacuum case:

Let  $\langle D, g, k \rangle$  be a vacuum data set. Then there is a unique, up to isometry, maximal vacuum Cauchy development (MVCD) of  $\langle D, g, k \rangle$ .

To explain, a vacuum Cauchy development of  $\langle D, g, k \rangle$  is a globally hyperbolic spacetime in which  $\langle D, g, k \rangle$  is embeddable. Two spacetimes  $(M, g)$  and  $(M', g')$  (here the metrics are Lorentzian) are isometric if there is a diffeomorphism  $\varphi : M \rightarrow M'$  such that  $\varphi^*(g') = g$ , where  $\varphi^*$  is a pull-back by  $\varphi$ .

The theorem does not prohibit a MVCD of an initial data set from having many non-isometric extensions—the theorem only prohibits these extensions from being globally hyperbolic. In this sense, the theorem does not offer support for determinism of GR, as it leaves it open whether there are non-isometric extensions of the MVCD for some vacuum data set. An example of such behavior is provided by the so-called polarized Gowdy spacetime. This is a globally hyperbolic spacetime and a vacuum solution to the Einstein equations. In accord with the theorem above, it has an MVCD of an initial data set. Yet, it is possible

to construct non-isometric extensions of the MVCD—see Chrusciel and Isenberg (Phys Rev D 48, 1993).

Moreover, Ringstrom (The Cauchy problem in general relativity, 2009) constructs two non-isometric extensions of MVCD for some specific initial data set (locally rotationally symmetric, Bianchi type IX), which (arguably) are extensions to the future. Such extension are naturally thought as alternative developments following a given vacuum data set, which is a paradigmatic picture of indeterminism. Given the symmetries inherent in the data set, it is not easy to dismiss this case as non-physical.

### **Matthias Egg. Views of the quantum state in Bohmian Mechanics and the GRW Theory.**

It has been noted for some time that Bohmian mechanics and the Ghirardi-Rimini-Weber (GRW) theory share a common structure (Allori et al., Brit. J. Phil. Sci. 59 (2008), 353-389). Both theories can be interpreted as postulating a 'primitive ontology' in space and describing its development in time. The difference between them lies in the kind of ontology that is postulated (particles in one case, a matter field or a set of discrete events in the other) and in its dynamics (deterministic and continuous versus stochastic and discontinuous). The aim of this paper is to further explore the commonalities and differences between these two theories, by analyzing the different roles that can be assigned to the quantum state within the two approaches.

For Bohmian mechanics, the ontological status of the quantum state has been discussed in recent papers by Belot (Eur. J. Phil. Sci. 2 (2012), 67-83) and Esfeld et al. (Brit. J. Phil. Sci. forthcoming, doi:10.1093/bjps/axt019). For the sake of comparing the Bohmian with the GRW context, it turns out to be helpful to group the different possible views of the quantum state according to how much autonomous existence they grant it. On one end of the spectrum, we then find the view that the quantum state is an object in its own right, understood either as a field on configuration space or some kind of generalized field, assigning properties to sets of points in ordinary space. On the other end of the spectrum is the Humean view that the quantum state has no ontological significance at all, but is merely a convenient tool to describe the temporal behaviour of the primitive ontology. As I will show, the discussion of these extreme views carries over straightforwardly from the Bohmian context to the GRW framework.

Things are different (and more interesting) for the positions closer to the middle of the spectrum. The view that the quantum state is a law (in a non-Humean sense, such that it actually governs the behaviour of the primitive ontology, rather than just describing it) is better motivated for Bohmian mechanics than for GRW, because the Bohmian picture, by virtue of introducing particle positions as additional variables not determined by the wave function, allows for a stationary quantum state of the universe even when there is movement in the primitive ontology, whereas this is not possible on the GRW view. Conversely, it can be argued that, if one regards the quantum state as a property of the primitive ontology, the matter density version of GRW (GRWm) leads to a more coherent ontology than the Bohmian framework. This is because, due to quantum entanglement, the quantum state must be a holistic property, which matches better with the holistic ontology of GRWm than with the atomistic one of Bohm. However, this raises the question as to the significance of 'particle' labels in the GRW wave function. In reply, I propose an interpretation of these labels in terms of the matter field's propensity for spontaneous localization.

**Lee Elkin. A Conciliation Model for Polarized Beliefs.** The Equal Weight View of peer disagreement is standardly modeled within a probabilistic framework where the difference between disagreeing parties is reconciled by averaging their degrees of belief and updating to the new subjective probability. The model seems appropriate for resolving fine-grained disagreements (e.g. disputes involving middling levels of confidence on some issue). But on various occasions, coarse-grained disputes arise that may be characterized as an ‘all-or-nothing’ affair. The standard model is unable to accommodate the latter kind of disputes for the reason that it is not at all clear on how to define qualitative or ‘all-or-nothing’ belief in probabilistic terms. Without having a precise definition of qualitative belief, it is unclear on how to split the difference between belief and disbelief within a numerical spectrum. To resolve the issue, I propose in this paper a different approach to splitting the difference between polarized beliefs that draws from the AGM theory of qualitative belief change. After setting up the formal framework, I illustrate how the AGM contraction operation, applied to each agent’s belief state, leads to splitting the difference where both arrive at a suspended state of judgment.

**Benjamin Eva. Interpreting Topos Quantum Theory.** It is now over a decade since Isham and Butterfield published the first papers on the topos-theoretic reformulation of quantum theory. In the intervening years, a remarkable amount of technical progress has been made in this area. The proponents of topos quantum theory (TQT) now routinely refer to the project as a ‘neo-realist’ formalisation of quantum mechanics. However, this development of an entirely new formalism for quantum mechanics, together with a fixed, ‘neo-realist’ interpretation that claims to overcome many of the traditional interpretational problems associated with the theory, has gone almost entirely unnoticed by philosophers.

In this talk, I will both provide an outline of the philosophical advantages that arise from taking TQT as the ‘correct’ formalisation of non-relativistic quantum theory (e.g. the fact that the logic that arises from the formalism is intuitionistic rather than orthomodular, the way that the Kochen-Specker theorem obtains a new, more intuitive, interpretation, the unification of logic and probability in the formalism, the fact that it is possible to simultaneously assign truth-values to all the physical propositions associated with a quantum system in a consistent way etc), and a critique of some of the interpretational difficulties that are still unresolved by the new formalism. Specifically, I will argue that the proponent of TQT does not have access to any single privileged notion of the physical state of a quantum system, and will consider the ways in which this difficulty relates to the purported ‘neo-realism’ of the theory.

Throughout, I will stress the philosophical analogies that hold between TQT and some of Bohr’s attitudes to quantum theory. In particular, I will argue that TQT can be seen as the natural formalism for anybody that takes Bohr’s ‘principle of complementarity’ seriously.

**Peter Fazekas, Balázs Gyenis, Gábor Hofer-Szabó & Gergely Kertész. A dynamical systems approach to causation.** Typically, philosophical approaches to causation follow one of two different routes. They either concentrate on providing an account of our everyday concept of causation as it features in causal discourse, or they try to capture what causation is in the objective world by uncovering the characteristic

features of causal relations as described by our best scientific theories. Counterfactual and difference making accounts famously follow the first route, whereas physical accounts of causation are more concerned about the second.

We find this duality unsatisfying. The fundamental aim of our paper is to invent a novel approach to causation that is able to integrate both objectives: to account for everyday causal claims and explain causal intuitions, on the one hand, and to do so in terms of how physical theories think about causal systems, on the other.

Causal systems are typically deterministic dynamical systems. Physical theories characterise deterministic dynamical systems by utilising two related concepts: a phase space, each point of which is a possible physical state that the system may find itself in, and a time evolution operator that describes how the physical states evolve with time.

Everyday causal discourse, on the other hand, when asserting causal claims like, for example, ‘in the presence of oxygen and combustible material, a short circuit causes fire’ relies on a set of natural linguistic descriptors (e.g. ‘there is a short-circuit’, ‘there is oxygen’) the referents of which are usually thought of as causes, effects, and background conditions.

Our framework is committed to the view that the properties picked out by natural linguistic descriptors supervene on physical states. Therefore, we identify the referents of natural linguistic descriptors with sets of physical states – and, thus, with regions of the phase space – for which the natural linguistic descriptors are true. Intersections of such phase space regions contain physical states that instantiate all the corresponding properties picked out by natural linguistic descriptors.

We propose that causal claims try to capture a relationship between different phase space regions. The phase space of any given causal system is such that there are regions in it from which all, or at least the overwhelming majority, of the physical states evolve into other regions. A causal claim is true if the properties postulated by the claim carve up the phase space in a way that they pick out these related regions.

That is, according to our approach, causal relations are phase space regions, the causal relation amounts to how the time evolution operator maps different phase space regions onto each other, and a cause of an effect (corresponding to an ‘effect-region’ in the phase space) can be any of those natural linguistic descriptors that together carve up the phase space such that they pick out a region related to the effect-region by the time evolution operator.

Moreover, our approach is able to solve most of the classical problems (such as causal selection, overdetermination, and causation by absences and misconnection) that pose serious difficulties for existing accounts of causation.

**Laura Felling. Causation, regularities and counterfactuals in fundamental physics: A solution to the bottoming-out problem.** The new mechanistic philosophy promises a simple solution to old issues in the philosophy of science: causation, counterfactuals and regularities. Such approach, however, breaks down at the level of fundamental physics. This is the so-called bottoming-out problem.

In this talk I consider the different aspects of the problem and put forward a solution.

In the first part of the talk I illustrate how the mechanistic view accounts for each of the above mentioned issues.

1) According to the mechanistic account of causation, causal relations are grounded on mechanisms. This account has

the virtue to allow for a singularist formulation, saving the intuition that singular causal relations can obtain even if they are not instances of a Law of Nature.

2) Counterfactual generalizations are justifiable with the knowledge of the mechanism underlying a behaviour, ‘without appealing to unanalysed notions of cause, propensity, possible world, or the like.’ (Glennan 1996 p. 63)

3) Regularities are accounted for by appealing to the robustness of mechanisms, rather than by positing a metaphysically thicker account of Laws of Nature.

Mechanisms are hierarchical: each interaction between parts of a mechanism also constitutes a mechanism. Under the assumption that such a regress cannot be infinite, though, there must be a bottom, i.e. the level of those interactions that are not underpinned by an underlying mechanism. Without such underlying mechanisms, the solutions prospected above break down.

In the second part of the talk I put forward a proposal to solve the bottoming out problem, in the three components that I have identified.

1) Causation. According to Stuart Glennan, if interactions at the fundamental level were not truly causal, then none of the putative higher-level causal relations mediated by mechanisms would be genuine (Glennan 2011). I defend the compatibility of the mechanistic account of causation with the claim that fundamental phenomena are non-causal. I argue that causation is an emergent feature of the world, appearing only within higher-level phenomena, where the behavior of complex systems can be explained in terms of the behavior of their components.

2) Counterfactuals. At the fundamental level, counterfactual claims are justified by the mathematical models displayed by the theory. Given that such models are also representations of the world, they allow us to perform surrogative reasoning, i.e. to translate the knowledge of the model into a knowledge of the world.

3) Regularities. Without robust mechanisms or Laws of Nature, how to explain fundamental regularities? The Humean way takes regularities as brute facts. I shall argue instead that regularities are not something that require an explanation. The quest for an explanation of regularities should therefore be rejected altogether. Its explanation is alien from the scope of a legitimate metaphysics of science.

**Samuel Fletcher. Global spacetime similarity.** A property of a scientific model is robust, or stable, when sufficiently similar models also have that property; a parameterized family of models varies continuously when small changes in the parameter accompany small changes in the model; and a sequence of models converges to another when certain relevant features of the models become arbitrarily similar to those of the limit model. These notions of similarity and continuity plays an important role in the analysis of scientific models and their interrelationships. It is not always obvious, however, how to make them precise.

In the context of spacetime theories, though, physicists have since the late 1960s employed tools from the mathematics of topology to encode precisely these kinds of relations: what it means for a one-parameter family of spacetimes to vary continuously, or for a property, such as obeying a causality condition or having a singularity, to be stable. The answers to these queries in specific cases will in general depend on which topology one chooses to place on the collection of all spacetimes. In practice, there have been two classes most often used, but insufficient at-

attention has been paid to examining the conditions under which these particular choices of appropriate. One exception is Geroch [1970, 1971], who has complained that these standard choices have undesirable features, presenting examples to show that one choice, the compact-open topology, seems to be too permissive regarding which sequences converge, while another, the open topology, seems to be too restrictive. I reconstruct Geroch’s examples and interpret them as illustrating a demand on the class of convergent sequences that a topology determines. In this light, his desiderata amount to a notion of uniform convergence or global similarity.

The principal result to be discussed is the construction of a topology satisfying his desiderata and the investigation of some of its properties. In particular, I show how the construction can be motivated physically, as corresponding to similarity of observations of certain classes of ideal observers, and mathematically, as respecting the (real vectorial) structure that the collection of spacetimes has. Finally, I remark on some possible applications, including the notions of approximate global spacetime symmetry and approximately conserved quantities.

### **James Fraser. Spontaneous symmetry breaking in finite systems.**

In both classical and quantum theories, systems with infinite degrees of freedom can have properties which are not found in any finite system. There has recently been some debate amongst philosophers of physics about the role these novel properties of infinite systems play in explaining and representing certain physical phenomena. A key ingredient of the orthodox approach to phase transitions in statistical mechanics is the thermodynamic limit. Roughly speaking, this means taking the model’s volume to infinity, which is evidently an idealisation of concrete systems that actually exhibit phase transitions. According to Batterman (2005) this idealisation is essential to the explanation of phase phenomena and has direct physical significance; he takes non-analyticities in the free energy found in the thermodynamic limit to correspond to genuine physical discontinuities that occur during a change of phase. This strong reading of the representational content of the thermodynamic limit has been contested by the likes of Butterfield (2011) and Callender and Menon (2013) however. These authors maintain that the thermodynamic limit has the more pedestrian function of bringing phase phenomena under mathematical control, while it is ultimately the de-idealised finite system which does the substantive explanatory and representational work.

This paper discusses a related phenomenon which raises similar questions about the status of the novel properties of infinite systems: namely spontaneous symmetry breaking (SSB). As with phase transitions, the most powerful approach to SSB in statistical mechanics seems to make indispensable use of properties which are only found in the limit of infinite degrees of freedom- in this case, the non-uniqueness of equilibrium states. I assess whether the deflationary reading of the thermodynamic limit, championed by Butterfield, Callender and Menon in the context of phase transitions, can be extended to SSB. The key issue, I suggest, is whether one can understand the thermodynamic limit as providing an appropriate representation of behaviour which is already present in large finite systems. I put forward two approaches to providing this kind of de-idealisation story. The first is based on an argument found in the classical statistical mechanics literature to the effect that, though the de-idealised model has a unique equilibrium state, we should expect it to enter and remain in an asymmetric state for very long time periods. The second follows recent work by Landsman (2013),

which demonstrates that the equilibrium state of quantum models which display SSB in the thermodynamic limit are unstable to asymmetric perturbations to the Hamiltonian. While the scope of these approaches is yet to be fully explored, I suggest that there is at least a program for deflating strong claims about the representational role of the thermodynamic limit in the context of SSB, though it does require some revision of textbook statements about what SSB is.

**Ellen Fridland. Intelligence automatically.** Despite being one of the most thoroughly explored phenomena in psychology over the last 50 years, philosophical notions of automaticity continue more or less unchanged, remaining largely simplistic and intuitive. This is not to say that psychologists themselves do not fall into simple dichotomies when thinking about automatic behaviors and processes but it is to say that philosophers seldom appreciate the complexity of automaticity. In this paper, I will review considerations in favor of a nuanced view of automaticity that avoids simple dichotomies. I will highlight a list of decomposable features, which are neither necessary nor sufficient for automaticity but that taken together, still form a theoretically useful cluster concept. After providing a broad overview of the psychological landscape regarding automaticity, I will highlight two substantive interactions between automatic and controlled processes such that classifying automatic processes as unintelligent becomes difficult. The first is the robust cognitive penetrability of automatic behaviors and processes and the second is the internal transformation of automatic behaviors and processes, which results from diachronic learning and practice. These considerations taken together should persuade philosophers to rethink their intuitive, simplistic notion of automaticity.

**Toby Friend. Laws as analysans for causation.** Laws of nature have often been used as analysans for token-level causation (see, e.g., Armstrong and Heathcote 1991, Armstrong 2004, Schaffer 2001, Maudlin 2004). However, the ways in which they have been used are highly variable, often seem ad hoc, and most importantly, pay no attention to the logical structure of laws as they are represented in science. For instance, the philosopher's caricature 'All Fs are Gs', and adaptations thereof, can seem of little use in representing either token causal relations (C caused E) or the complex (and often *ceteris paribus*) dynamical equations and qualitative relations called 'laws' (e.g. the Lotka-Volterra equations, Pauli's Exclusion principle).

Given these failures are for such distinct reasons, one might be sceptical of an analysis of causation either in terms of the philosophical caricatures or scientific renderings of laws. It is, therefore, surprising to find how often laws are employed to shore up causal theories, (e.g. in Lewis 1973, Dowe 2000, Paul 2000, Hall 2007). Indeed, Tim Maudlin has claimed that we need laws as analysans of causation, but admits, 'I do not think that there is any uniform way that laws enter into the truth conditions for causal claims' (2004, 430). I agree with Maudlin's claim, but disagree with his scepticism; my purpose in this paper will be to attempt a uniform way.

My approach begins, in 1, with an examination of past attempts to analyse causation in terms of laws. I argue, drawing on electronics examples, that these attempts fail because of a misconception of the logical form of (non-causal) laws relevant to causal analysis. In 2, I develop an understanding of how 'All Fs are Gs' can display the logical form of electronics laws. Roughly,

an assembly of electronic components can be expected, according to some law, to behave in a way specified by the law's consequent 'complex behavioural property' G (e.g.  $V=IR$ ) when it satisfies characteristics specified by the law's antecedent 'instantiation conditions' F (e.g. being a closed loop of conductive material). An assembly satisfying a (group of) law's instantiation conditions is called a system and, in 3, I introduce a basic kind of link between objects, the Intra-system link, showing how some simple causal relations (e.g. a resistor's variation causing a bulb to illuminate in a single circuit) can be characterised in its terms. In 4, I consider more complex forms of link between objects, Inter-system links and System-corruption links, and show how these can characterise further instances of causation (e.g. respectively, a switch closing causing a bulb to illuminate, and a switch opening in one circuit causing a bulb in another 'competing' circuit to illuminate). Causation is thus defined as a chain of combinations of system-links of the sorts described. I end in 5, by briefly considering the account's ability to generalise and its merits regarding treatment of pre-emption and negative causation.

**David Glass & Mark McCartney. Explanatory competition and explaining away.** What does it mean for two hypotheses to be in competition with each other? The answer to this question is relevant in the context of debates about evidential favouring, inference to the best explanation and explaining away arguments where reasons to accept one explanation are put forward as reasons to reject another explanation. In some discussions, the only type of competition considered is that which occurs between mutually exclusive hypotheses, but it seems clear that there are cases where compatible hypotheses can compete with each other. This topic will be explored in the context of explanatory hypotheses.

One case in which two compatible explanatory hypotheses might be considered to compete is when there is a negative probabilistic dependence between them, but even if they are probabilistically independent, competition can still occur. To take a simple example, suppose that my car will not start and two possible explanations spring to mind: a flat battery and an engine problem. Even though these two explanations are independent of each other, when it is discovered that the battery is flat, this counts against the alternative explanation that there is an engine problem; there is no need to infer two explanations when one will do. The reason that such 'explaining away' occurs is that although the explanations are unconditionally independent, they become conditionally dependent on the evidence that the car will not start.

This 'explaining away' mechanism occurs frequently in probabilistic networks which have been a subject of considerable interest in Artificial Intelligence over the last thirty years. It is important to note that explaining away does not always occur when there are two possible explanations of a piece of evidence. In some cases the explanations can mutually enhance each other and so explanations which are negatively dependent before the evidence is taken into account will not necessarily still be negatively dependent afterwards.

This paper explores these issues using Bayesian confirmation theory in order to construct an account of explanatory competition. The potential relevance of this work in several areas of philosophy of science will also be considered briefly.

**Teddy Groves. Accuracy arguments in the context of Carnapian inductive logic.** Several recent philosophical arguments seek to show that states of belief must be repre-

sentable by probability spaces in order to avoid being needlessly inaccurate. I consider whether such accuracy arguments can be applied to the project of developing Carnapian inductive logic by supporting the claim I call ‘probabilistic necessity’. I argue that they cannot.

I begin by introducing inductive logic. An inductive logic is a pair  $(L, m)$  consisting of a formal language  $L$  and a ‘measure function’  $m$  that associates  $L$ ’s sentences with real numbers. The Carnapian tradition in inductive logic investigates ‘adequacy criteria’: rules labelling certain inductive logics impermissible. Collections of such adequacy criteria, it is hoped, can usefully replace informal inductive assumptions.

A key issue in Carnapian inductive logic concerns the status of probabilism. Probabilism is the adequacy criterion that permits only inductive logics whose measure functions are probability functions. In particular, it is controversial whether probabilism should be a part of all collections of adequacy criteria. I call this claim ‘probabilistic necessity’.

Probabilistic necessity can be supported by epistemological arguments to the effect that all rational states of belief are representable by probability spaces. Accuracy arguments for this claim have been advanced in the recent formal epistemological literature by Joyce, Predd et al. and Leitgeb and Pettigrew.

In general, such arguments consist of assumptions about rational states of belief—that they concern sentences of a propositional language, are representable by real-valued belief functions and are not dominated with respect to accuracy by other states of belief—together with descriptive claims about the nature of inaccuracy.

Although I note that each of the accuracy arguments’ propositions about states of belief is questionable, the main aim of my talk is to contest some prominent claims about the nature of inaccuracy, namely that legitimate measures of inaccuracy are sum decomposable, strictly proper and continuous.

A sum-decomposable measure makes each state of belief’s inaccuracy the sum of its proposition-specific inaccuracies. I argue that, although some form of relationship between proposition-specific and global inaccuracy is plausible, there is no compelling reason why it should be expressed by summation, rather than another operation.

Sum-decomposable inaccuracy measures are strictly proper if, for all propositions  $X$ , and probabilities  $p$ , an agent can only minimise expected inaccuracy with respect to  $p$  and  $X$  by setting their belief  $b(X)$  equal to  $p$ . I argue that strict propriety is difficult to justify without first assuming that rational agents have probabilistic states of beliefs.

Continuous inaccuracy measures are sum-decomposable and additionally are such that, for all propositions  $X$ , small changes in degree of belief in  $X$  lead to small changes in proposition-specific inaccuracy with respect to  $X$ . I argue that this condition unjustifiably rules out some potentially plausible inaccuracy measures.

In light of the difficulties I find with assuming that inaccuracy has these properties, I conclude that the prospects for using accuracy arguments to justify probabilistic necessity are bleak.

**Andreas Hüttemann. Actual causation and default processes.** It has been suggested that a theory of causation is in need of characterizing certain kinds of behaviour as default behaviour (Hall 2007, Halperin and Hitchcock, forthcoming). In this paper I provide an account of actual causation that takes default processes as its central notion.

The first step of my analysis consists in analysing actual causation in terms of an interfering factor to a default process: A cause is an interfering factor with respect to the default behaviour/ process of a system. (This covers many but not all cases of actual causation.)

The second step consists in further analysing default-processes and interfering factors in scientific terms, in order to yield an account that is reductive in the following sense: causal facts are shown to be nothing over and above science facts (empting those facts e.g. in the social sciences that are explicitly characterized in causal terms). Traditional process-theories provide a one-size-fits-all characterization in terms of the transmission of a certain amount of a conserved quantity. This led to a number of problems (see below). What is essential is (1) that it is a determinate and objective fact what the default-process of a system is and (2) that the default-behaviour can be characterized in scientific terms: Default processes are those processes that systems are disposed to display provided there are no interfering factors. Newton’s first law describes the (quasi)-inertial or default behaviour of a massive particle. They are disposed to display a certain behaviour (‘continues in its state of rest or of uniform motion in a straight line’) provided there are no interfering factors (‘less it is compelled to change that state by forces impressed upon it.’). ‘Laws of deviation’ (Maudlin 2004, 431) such as Newton’s second law determine how interfering factors obviate the default behaviour. It is argued that in other sciences as well as in ordinary contexts we have analogous information about default processes and interfering factors. Default-processes and interfering factors are identified and characterized in the sciences (and elsewhere) on a case-by-case basis.

This account can deal with the major problems that have been raised for the conserved quantity theory: (1) It is not committed to an implausible form of reductionism: It does not require all talk of causation to be translatable into talk about conserved physical quantities. (2) Disconnections: Negative causation cannot be integrated into process theories that require the persistence of physical characteristics along a world-line that connects cause and effect. According to the account presented here, the absence of a feature of a system may be due to an interference into a default-process. (3) Misconnections: Only some interactions seem to be causally relevant. The conserved quantity theory cannot clarify this distinction. On my account a factor interferes with a default process if – as a matter of fact – the factor obviates the display of the default behaviour. ‘Interfering factor’ is a success term. What it means that a factor not only to interact but to interfere can be spelt out in terms of ‘laws of deviation’.

**Tero Ijäs. Beyond tinkering: Design and understanding through directed evolution in synthetic biology – a case from protein design.** Synthetic biology is a growing post-genomic research field which aims to construct artificial biological components and systems. One of its core principles is argued to be the application of engineering-inspired ‘rational design’. Rational design uses computer simulations and fabrication to reduce complexity and create well-defined standardized synthetic components and devices. Ideally, these simple parts would be modular enough to be swapped or combined, and could be assembled together to create more complex genetic devices or systems with novel functions. However, in many cases the created devices are sensitive to cellular context and their behavior is susceptible to unpredictable changes when these devices are implemented to different system or combined with



other components.

Therefore, in more complex cases rational design has been complemented or replaced with other design methods, such as 'directed evolution'. In directed evolution researchers start with the selection of mutational target that they hope to modify and induce mutational change in it. The aim is to generate a library of DNA molecules with varying phenotypes. Desired phenotypes are screened and selected from the generated library, and – through multiple iterations – guided towards desired circuit functionality. Directed evolution allows exploration of different options in engineering of synthetic circuits in the cases where researchers lack enough structural information of used systems or when the iteration between modelling, redesign and construction has not provided working circuits. Even though the use of directed evolution and other similar methods is widespread in synthetic biology, closer philosophical analysis of their role and heuristics has so far been scarce. Directed evolution has in many instances been categorized as an 'irrational' design approach which prioritizes practical question of circuit functionality and optimization over understanding. O'Malley (2011) sees directed evolution as a method of 'kludging', a way of creating particular, pragmatic solutions to design problems that is closer to 'tweaking' or 'debugging'.

In this paper, I elaborate the role that directed evolution plays in synthetic design by analyzing its application in protein design. Redesign and construction of proteins with optimized or novel functionality is seen as a major goal in synthetic biology. However, due to combinatorial complexity of possible protein structures, rational design approach is found lacking and protein design is carried mainly out by directed evolution. I will analyze design choices that the use of directed evolution requires in different stages of protein design (e.g. choice of redesigned parent protein and mutational target sequence). I will also argue why claimed trade-off between functionality and understanding is not a necessary consequence of directed evolution, but depends from the non-modularity of the designed system and how its coupling with the environment changes. Finally, I will argue that protein design offers a case where directed evolution is applied systematically with proper understanding of its requirements and limitations, and its results more 'irrational' or 'kludge-like' than those of rational design.

### **Elselijn Kingma. Metaphysics of pregnancy: Fetuses as part of the maternal organism.**

I take these two statements to be uncontroversial: (1) before an organism becomes pregnant, it is only one organism. 2) after the organism's pregnancy, there are (usually) two organisms. Together, these two statements raise a question: when does one organism become two?

Smith & Brogaerd (2003) provide an answer: 16 days after (human) conception. This, they argue, is when gastrulation starts; when the human embryo changes from a clump of cells to a three-dimensional, differentiated, multicellular whole; and when that embryo first becomes the same organism as the future baby and adult organism that it will grow into. Moreover, Smith & Brogaerd argue, the embryo's location inside the maternal womb poses no problem for transtemporal identity; in line with our popular depiction of pregnancy, they argue that fetuses are not part of a pregnant organism, but merely inside that organism – like a tub of yogurt is in the fridge, or a 'bun is in the oven'.

It is this latter claim that I contest in this paper. Using Smith & Brogaerd's own criteria for topological connectedness, and ex-

amining three different conceptions of the foetus, I argue that there is no physical discontinuity between foetus and maternal organism; whether we look at the umbilical cord or at the placenta, foster and mother are firmly topologically connected, as well as functionally and metabolically integrated. Thus, I argue, fetuses are a part of the maternal organism – just as much as (her) kidneys, blood or hair are – up until birth.

Smith & Brogaerd (and many others, e.g. Olson, 1997) take us to be organisms. If we follow Smith & Brogaerd's conception of the organism, where organisms are numerically and physically distinct substances, this means that we come into existence at birth, and no earlier. This view has many advantages: it is numerically neat; it aligns with an intuitive picture of organisms on which organisms are always distinct individuals; and it ties coming into existence to a distinct event, birth, thus preempting the need to draw dichotomous boundaries during the previous nine-month period. Nevertheless it has a distinct drawback: we were never fetuses: birth is, on this view a substantial change, and no foetus could ever survive it.

Of course we need not accept that conclusion; we could either reject the view that we are organisms, or we could reject Smith & Brogaerd's view of organisms, such that they can be part of other organisms of the same kind. I only discuss the latter option. Whilst rejecting the substance-view of organism is probably correct, this still has a significant cost on any view of humans as organisms: first, it leaves the question of when an organism comes into existence again unanswered. Second, it means that human organisms are quite different entities than we thought they were: they can be part of each other.

### **Bon-Hyuk Koo. How much can we grasp? Objective blind realism as an answer to pessimistic meta-induction and Stanford's 'trust' argument.**

In this paper I present a new viable and minimal form of scientific realism, which is named 'objective blind realism' that derives from a similar position asserted by Robert Almeder (1987, 1994).

Pessimistic meta-induction (PMI), a major anti-realist argument, capitalises on theoretical discontinuities across theory change to establish falsity of past successful theories, and then to infer that our theories are likely to follow suit and/or that success of scientific theories does not guarantee approximate truth. The realist responses have tried to show that there are also continuities across theory change and that those theoretical parts are approximately true. Various realist positions point at different parts of scientific theories ranging from theoretical parts that indispensably contribute to predictive successes (Psillos 1999) to entities (Hacking 1983, Cartwright 1983), structure (Worrall 1989, Ladyman 1998), and properties (Chakravartty 1998, 2007) to mention a few.

Without taking side in a particular realist position, I put forward objective blind realism (OBR). Almeder's blind realism and OBR are similar in holding that we can trust certain parts of our scientific theories to be approximately true descriptions of nature even if we do not know which parts they are. The differences are that (1) unlike Almeder's, OBR does not rule out the possibility of identifying approximately true theoretical parts; and (2) OBR endorses a correspondence theory of truth whereas Almeder's is based on a coherence theory of truth. Success of scientific theories gives us reason to believe in their being approximately true by virtue of the No Miracle Argument (NMA). As long as NMA stands, success of a scientific theory can be attributed to the theory corresponding with the reality, contra Almeder's coherentist blind realism.

When faced with PML, OBR shares with the other realist accounts the necessity of finding continuity across theory change. However, its advantage over other positions is that OBR need not go any further to establish that the retained theoretical constituents are indeed approximately true parts. OBR avoids having to identify the exact parts of theory to place our epistemic commitment, while not falling into anti-realism.

This advantage is made prominent by Stanford's 'trust' argument (2003, 2006) that moves anti-realist scepticism from the level of scientific theories to that of scientists by pointing out their repeated failures to place trust in the 'wrong' parts of scientific theories, discrediting the attempts to delineate theoretical parts worthy of the realist commitment – illustrating scientists and philosophers, past and current, as unreliable epistemic agents. The crux of the debate with Stanford's argument, as well as between different realist accounts, then becomes whether there is a tenable set of prospective criteria for identifying which parts of theories are likely to be retained, and if so, what they may be. I examine two notable realist attempts (Psillos 1999 and Chakravartty 2008) and show them to be wanting. While OBR may be seen as too weak even for realists, the position is significant in that letting go of the requirement for such prospective criteria still allows justified optimism for realism.

### **Katharina Kraus. Quantifying introspection? – The case of pain measurement.**

Psychologists often dismiss introspection as an inappropriate research method for various reasons: the data gained through introspection are highly inaccurate and cannot be objectively justified or numerically quantified (Schwitzgebel 2008). Yet a subject's report on her own experience often contains information that is highly valuable for understanding the underlying psychological processes. This paper will explore the possibility of quantifying first-person experiences (expressed in self-reports) and their correlatability with third-person data acquired from behavioural experiments and physiological measurements. This issue will be discussed in the context of pain measurement. It will be argued that valid quantitative data can be obtained through introspection, though only if an appropriate conception of psychological measurement is presupposed.

Pain is commonly viewed as a personal, subjective experience influenced by cultural learning, the meaning of the situation, attention, and other psychological variables. Approaches to the measurement of pain include self-rating scales, behavioural observation scales, and physiological responses. It has repeatedly been argued that, due to the subjectivity of pain, self-report based measures are the most valid and accurate tools for pain measurement (Katz et al. 1999; Noble et al. 2005). Nevertheless, there are ongoing attempts to find purely physiology-based measurements (Brown et al. 2011).

Philosophers of science have doubted whether the psychological attributes observed in such measurements are in principle quantifiable in any meaningful sense. Joel Michell recently raised new concerns and argued for the so-called quantity objection in psychology (e.g., Michell 1999; 2006). By discussing Michell's arguments with respect to pain measurement, I will show that there are good reasons to reject the quantity objection and to accept measurements based on subjective scales of intensity. I agree, however, with Michell that, in order to justify the measurability of such psychological attributes, the classical conception of measurement, according to which all measurable attributes are quantitative, has to be substantially amended. Following Steven's operationalist interpretation of psychologi-

cal measurement (Stevens 1951), I will contend that measuring processes should be viewed as conventional ordering activities that are historically and socially situated. Yet I will modify this operationalist approach insofar as I will argue that these ordering activities are governed by certain conceptual constraints that are due to the nature of our mathematical-quantitative reasoning, rather than due to the nature of the measured phenomena. I will conclude that any theory of quantitative psychology will have to take these conceptual constraints into account.

Brown, Justin et al. (2011) 'Towards a Physiology-Based Measure of Pain'. *PLoS One* 6(9):1-8.

Katz, J. Melzack, R. (1999) 'Measurement of Pain'. *Surg. Clin. North Am.* 79(2):231-52.

Michell, Joel (1999) *Measurement in psychology*. Cambridge: CUP.

Michell, Joel (2006) 'Psychophysics, intensive magnitudes, and the psychometricians' fallacy'. *Studies in History and Philosophy of Biological and Biomedical Sciences* 17:414-432.

Noble, Bill et al. (2005) 'The Measurement of Pain, 1945–2000'. *Journal of Pain and Symptom Management* 29:14-21.

Schwitzgebel, Eric (2008) 'The unreliability of naive introspection'. *Philosophical Review* 117:245-273.

Stevens, Stanley (1951) 'Mathematics, measurement and psychophysics'. In: S. Stevens (ed.) *Handbook of experimental psychology*.

### **Jaakko Kuorikoski & Caterina Marchionni. Evidential diversity and the triangulation of phenomena.**

The paper clarifies the epistemic rationale of triangulation as a form of robustness analysis (Wimsatt 1981), that is, as the use of multiple and independent sources of evidence to ascertain whether a phenomenon is an artifact of a particular method. The notion of triangulation as robustness analysis is closely related to that of evidential diversity, but although the confirmational significance of evidential diversity is a widely accepted epistemic principle (e.g. Fitelson 2001), several worries about robustness analysis have been voiced (e.g. Stegenga 2009, Hudson 2013). For example, in a challenging critique, Jacob Stegenga (2009) has recently argued that robustness analysis faces several difficulties that limit its epistemic value, namely, that evidence produced with different methods is often incomparable, that a criterion of independence is needed but is not available, that robustness analysis does not always work as a confirmatory procedure, and that multiple methods often yield results that are not congruent.

We defend triangulation as a form of robustness analysis against these challenges. We show that in order to evaluate its epistemic benefits, two kinds of inferences should be distinguished: inferences from data to phenomena and inferences from phenomena to theories (Bogen and Woodward 1988). Triangulation does not work in the same way in the two cases and the requirements for inferring the robustness of a result are different from those required in bringing a variety of evidence to bear on a theory. Unlike theory-phenomena inferences, data-phenomena inferences concern the causal processes generating the data (Bogen and Woodward 1988). In robustness arguments about phenomena, what we need to worry about are errors in the particular processes at play.

On this account, triangulation is to be understood as an epistemological strategy employed for the purpose of controlling for likely errors and biases. Independent experimental procedures or kinds of evidence can be used to increase the 'aggre-

gate' reliability of the evidence for a phenomenon. The relevant notion of independence is error independence, which does not require knowledge of all the problematic background assumptions of different methods as seems to be required for the confirmational boost that a variety of evidence confers to a theory. Furthermore, error independence is established case by case and does not pose pressing issues of incomparability of evidence; what is needed instead is (preferably controlled) co-variation between the phenomenon and the data. Our main illustrative case will be social scientific experiments on cooperative behavior and social preferences.

Bogen, J. and J. Wooward (1988) 'Saving the Phenomena.' *The Philosophical Review* 97: 303-352

Fitelson, B. (2001) 'A Bayesian account of independent evidence with applications.' *Philosophy of Science* 68: S123-140

Hudson, R. (2013) *Seeing Things. The Philosophy of Reliable Observation.* Oxford University Press

Stegenga, J. (2009) 'Robustness, discordance, and relevance.' *Philosophy of Science* 76: 650-661

Wimsatt, W. (1981) 'Robustness, reliability, and overdetermination.' In M. Brewer and B. Collins (eds.) *Scientific Inquiry in the Social Sciences*, San Francisco: Jossey-Bass: 123-162

**Jürgen Landes. Strictly proper scoring rules and the Probability Norm.** One of the most widely accepted norms of rational belief formation is the Probability Norm which requires agents to adopt beliefs that satisfy the axioms of probability. For example, the Probability Norm is held dear by all Bayesians. The question arises as to how to justify this norm. Traditionally, axiomatic justifications and Dutch Book Arguments were given to this end. The latter have been widely regarded as the most persuasive justification, however they have recently begun losing some of their once widespread appeal [Hájek, 2008].

Recent work in epistemology takes a less pragmatic approach using epistemic scoring rules to justify the probability norm [Joyce, 1998, Joyce, 2009, Leitgeb and Pettigrew, 2010, Predd et al., 2009]. Scoring rules were first studied by Brier in 1950 as a tool to elicit probabilistic degrees of beliefs from forecasters. Brier's work has been highly influential in the statistical community which developed the notion of a statistical scoring rule, which made its way in the *Encyclopedia of Statistics*, see [Dawid, 1986, p. 211]. Epistemic scoring rules differ in form and application from statistical scoring rules.

I will argue that statistical scoring rules, properly understood, are in principle better suited than epistemic scoring rules to justify the Probability Norm. My argument proceeds as follows: The most convincing justifications of the Probability Norm relying on epistemic scoring rules require the scoring rules to have a certain property, strict propriety. However, for purposes of justifying the probability norm, assuming that an epistemic scoring rule is strictly proper is question begging. On the contrary, strict propriety for statistical scoring rule is not only defensible but a desideratum.

The mere argument that statistical scoring rules are in principle better suited to justify the Probability Norm does not get us closer to a convincing justification of the Probability Norm. In the second part of this talk I will note how to use statistical scoring rules to justify the Probability Norm. I will also touch on how to use statistical scoring rules to justify Maximum Entropy Principles and a probabilistic Principle of Indifference.

Dawid, A. P. (1986). Probability forecasting. In Kotz, S. and Johnson, N. L., editors, *Encyclopedia of Statistical Sciences*, volume 7, pages 210–218. Wiley.

Hájek, A. (2008). Arguments for – or against – Probabilism? *British Journal for the Philosophy of Science*, 59(4):793–819.

Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4):575–603.

Joyce, J. M. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*, volume 342 of *Synthese Library*, pages 263–297. Springer.

Leitgeb, H. and Pettigrew, R. (2010). An Objective Justification of Bayesianism. *Philosophy of Science*, 77(2):201–272.

Predd, J., Seiringer, R., Lieb, E., Osherson, D., Poor, H., and Kulkarni, S. (2009). Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions on Information Theory*, 55(10):4786–4792.

**Tim Lewens. The perils of cultural models.** A great deal of work done under the banner of cultural evolutionary theory involves the construction of idealised explanatory models: indeed, the use of such models is often touted as the key virtue of the cultural evolutionary approach. It is important, therefore, to understand both generic worries one might have about the application of such models to the cultural domain, and also specific worries about the construction of particular cultural models. I begin by demonstrating the ubiquity of model building in cultural evolutionary studies, before considering a pair of very general accusations brought against cultural modelling by Tim Ingold. Ingold believes that cultural evolutionary theorists use models in a manner that is circular, and which thereby offers only spurious confirmation to their proposed explanatory hypotheses. He also believes that the processing of ethnographic data that is required to make them apt for modelling purposes—more specifically the manner in which they must be abstracted from the context in which they were gathered in order to render them suitable for mathematical formalisation—undermines the reliability of these data.

Ingold's criticisms, at least when read at face value, both fail. This does not mean that cultural modelling is in the clear. Cultural evolutionary models frequently offer dubious formalisations of the hypotheses they claim to test. Moreover, the support the models' assumptions receive from direct experimental work is often exaggerated. Close attention to specific examples of cultural model-building shows that concerns about circularity and the use of data away from the context of their generation do offer sources of legitimate concern, although not for the reasons Ingold envisages. While both criticisms lead on to suggestions for how to use cultural models more persuasively, neither criticism entails that cultural evolutionary modelling should be abandoned: there is no other means by which we can explore hypotheses about the manner in which populational cultural patterns are produced by the aggregated effects of individual interactions.

**Chiara Liscandra. Robustness Analysis as a Non-Empirical Confirmatory Practice.** Robustness analysis is a method of testing whether the predictions of a model are the unintended effect of the unrealistic assumptions of the model. As such, the method resembles the analysis, conducted in experimental sciences, to test the effect of possible confounders on the empirical results. The arguments in support of

robustness analysis in non-experimental contexts, however, are often left implicit or are unreflectively imported from the experimental sciences. The aim of this paper is to cast light on the logic behind robustness analysis and to examine the criteria in its support.

More specifically, this paper focuses on the problem of robustness with respect to tractability assumptions, i.e. different mathematical formulations of the same factor in a model. I will show some difficulties this method encounters in scientific practice and argue that the very possibility of conducting tractability robustness analysis deserves further clarification. If robustness analysis were a ‘surgical’ operation, in which controversial aspects could be replaced by other ones with no other relevant changes, then the role of a single assumption could be evaluated and the consistency of the results after variation assessed. Yet, it is not always possible to introduce changes in a model without altering its main structure. It is more often the case that the intimate connection between simplifying assumptions and mathematical tractability is such that variations can only be introduced by altering the overall structure of the model, which eliminates the possibility of analyzing the effect of one specific change.

It has been urged by several authors (Cartwright 2006, Kuorikoski 2010), that the impact of tractability assumptions requires a systematic analysis, which is still missing in the philosophy of economics literature. The present paper aims to move forward in this direction. It elucidates the rationale underlying robustness analysis, highlights the difficulties encountered in the practice, and indicates where effective strategies need to be developed in response to these difficulties.

**Sebastian Lutz. Abstraction, idealization, and the application of mathematics.** I explicate the notions of ‘idealization’ and ‘abstraction’ as special instances of distortions and omissions, respectively, so that they are inferentially relevant. The results clarify the applicability of mathematics in science.

To distort a description (a set of sentences)  $\Theta$  is to give a description  $\Delta$  that is incompatible with  $\Theta$  given some fixed set of background assumptions  $\Lambda$ . A description  $\Delta$  thus distorts a description  $\Theta$  only if  $\Theta \cup \Delta \cup \Lambda \models \perp$ . A description  $\Omega$  omits from a description  $\Theta$  if and only if  $\Theta \cup \Lambda \models \Omega$  and  $\Omega \cup \Lambda \not\models \Theta$ . I give a simple example of an irregular quantum well being distorted into a rectangular quantum well. The descriptions and its distortion have a nontrivial common omitting description of the wave function.

I suggest treating abstractions as omissions of all sentences that contain a specific vocabulary. For a set  $\theta$  of sentences in vocabulary  $V$  and  $S \subseteq V$ , define the  $S$ -consequences of  $\Theta$  as  $\Theta|S := \{\sigma : \Theta \models \sigma \text{ and } \sigma \text{ is an } S\text{-sentence}\}$ . Call a description  $A$  an abstraction of a description  $\Theta$  in terms of  $S$  if and only if  $A$  omits all and only those sentences that cannot be inferred from  $\Theta$ 's  $S$ -consequences:  $A \cup \Lambda \models (\Theta|S) \cup \Lambda$ . I discuss this for the quantum well example. Idealizations are commonly accepted to be distortions. If abstractions are supposed to justify idealizations in the way that omissions justify distortions, a description and its idealization should have a common abstraction. Then a description  $I$  idealizes a description  $\Theta$  in the vocabulary  $S$  only if it distorts only consequences of  $\Theta$  that contain terms not in  $S$ :  $\Theta \cup I \cup \Lambda \models \perp$  and  $(I \cup \Lambda)|S \models (\Theta \cup \Lambda)|S$ .

Pincock (2007: ‘A Role for Mathematics in the Physical Sciences’, *Noûs* 41, 255) argues that ‘mathematics allows us to make claims about higher-order or large-scale features of physical systems while remaining neutral about the basic or micro-

scale features of such systems’. He uses Euler’s solution to the Königsberg bridge problem to argue that when scientists accept a mathematical statement like ‘The bridge system forms a non-Eulerian graph’, they implicitly ignore those mathematical properties they believe are inappropriate to ascribe to the physical system (263). Hence Euler’s proof does not rely on the microstructure of Königsberg, and so does not ‘fail if the microphysics of the bridges [is] altered’ (260). But Pincock’s reconstruction of Euler’s proof does rely on a (false) microstructure of Königsberg, namely a definition of ‘vertex’ and ‘edge’ in set theoretic terms. The proof is robust only if set theory is ignored. But it is not at all clear why or which parts of mathematics can be ignored, and thus it is not clear why or whether Euler’s proof is robust. I will argue that moving from the definition in set theoretic terms to a set of axioms for ‘vertex’ and ‘edge’ shows how mathematics allows us to make claims about large-scale features of Königsberg. The axioms are abstractions of both the reductive explication and of the bridges of Königsberg. Hence it is abstraction, not mathematics, that leads to robustness and the applicability of mathematics in science.

**Daniel Malinsky. Hypothesis testing, ‘Dutch Book’ arguments, and risk.**

‘Dutch Book’ arguments and references to gambling theorems are typical in the debate between Bayesians and scientists committed to ‘classical’ statistical methods. These arguments have rarely convinced non-Bayesian scientists to abandon successfully applied and ubiquitous ‘classical’ methods, such as null hypothesis significance testing, partially because many scientists feel that gambling theorems have little relevance to their research activities (testing psychological response theories, mapping neural structures, looking for novel particle phenomena, etc.). In other words, scientists ‘don’t bet’. This paper examines one attempt, by Schervish, Seidenfeld, and Kadane, to progress beyond such apparent stalemates by connecting ‘Dutch Book’-type mathematical results with commonly endorsed ‘classical’ statistical principles. A theorem that SSK prove in their (2002) paper seems to lay the foundation for a normative argument against the scientist committed to testing simple null hypotheses at a conventional (fixed) alpha level, such as  $\alpha = .05$ . I argue that such a normative argument would fail to be convincing to any non-Bayesian scientist, because their mathematical result fails to be interpretable in the light of experimental practice. Luckily, SSK’s work does suggest a way to move the debate forward. By focusing the conversation on a statistical procedure’s associated risk and the scientist’s preference for minimal risk functions, the Bayesian criticism can point to conflicts between some principles a ‘classical’ experimentalist does actually endorse. In particular, testing hypotheses at a fixed alpha-level conflicts with the commonly espoused preference for minimal risk functions, but only when that preference is extended to cover the ‘average’ or ‘combined’ performance of a series of tests. If a scientist is interested in minimizing the overall risk of their testing procedure, they have good reasons to adopt a coherent rejection rule, i.e., they ought to act as if their prior probabilities over their hypothesis space satisfy the requirements of coherence. Furthermore, if an experimentalist adopts a coherent rejection rule, they can do better in terms of minimizing Type I error probability. I conclude with some comments on contemporary experimental particle physics, where something like a conventional fixed alpha-level rejection rule for particle discovery is currently enforced by journals (the 5-sigma rule). Experimental physicists, who often emphasize the importance of minimizing Type I error probabilities, can do better by their own lights if

they adopt a coherent testing procedure.

Schervish, M.J., Seidenfeld, T., and Kadane, J.B. (2002), 'A Rate of Incoherence Applied to Fixed-Level Testing', *Philosophy of Science*, Vol. 69, No. S3: S248-S264.

### **Alexandru Marcoci. Solving the absentminded driver problem through deliberation.**

Piccione and Rubinstein [1997] have suggested a sequential decision problem with absentmindedness in which there seem to be two equally compelling, but divergent, routes to calculating the expected utility of an agent's actions. The first route would correspond to an ex ante calculation while the latter to an ex interim calculation. Piccione and Rubinstein conclude that since the verdicts of the two calculations disagree they lead to an inconsistency in rational decision theory.

In this paper I firstly argue that the ex ante route to calculating expected utility is not available in decision problems such as that introduced by Piccione and Rubinstein. The reason is that ex ante expected utility requires the agent to have a vantage point before making any decision from which to contemplate all the possible paths through the decision tree. This is reasonable to expect in games without absentmindedness since the root of a tree without absentmindedness cannot be part of a non-singleton information set. However, in the absentminded driver problem, the root of the tree, which is also the first decision node for the driver, is included in the same information set as the second decision node for the driver. Therefore, there is no vantage point from which the driver can contemplate all the possible complete paths through the tree: when the driver is at the first decision node, for all he knows, he may be at the second one. And there, the ex ante expected utility formula does not make sense.

The second part of the paper will explore the ex interim expected utility formula. This has been largely neglected in the literature and is always presented as only offering the driver a parametric optimal strategy in terms of his initial belief in being at the first decision node. I will argue that if we construe agents as maximising the ex interim expected utility in steps through a deliberative dynamics, then this formula can make a precise recommendation with regards to the driver's optimal strategy irrespective of his initial beliefs in being at the first decision node.

### **Carlo Martini. The limits of trust in interdisciplinary science.**

In this paper I argue that lack of trust networks among researchers employed in interdisciplinary collaborations is potentially hampering successful interdisciplinary research. I use Hardwig's concept of epistemic dependence in order to explore the problem theoretically, and MacLeod and Nersessian's ethnographic studies in order to illustrate the problem from the viewpoint of concrete interdisciplinary science practice. I suggest that some possible solutions to the problem are in need of further exploration.

The topic of pluralism in science has long been linked with the topic of interdisciplinarity: 'The appreciation of the need for interdisciplinary approaches in science studies aligns with pluralism at the metaphilosophical level. Because the scientific enterprise is itself a complicated phenomenon, no single disciplinary approach can provide a fully adequate account of its conceptual, technical, cognitive-psychological, social, historical, and normative aspects [...]' (Kellert, Longino and Waters 2006, ix) For similar remarks see also Dupré (2001).

Despite the importance of the topic, the several articles that have been written on interdisciplinarity, an edited volume

(Frodeman, Klein and Mitcham 2010), and a special issue of the journal *Synthese* (VV.AA. 2013), interdisciplinarity tends to be an elusive subject in philosophy. The problems and challenges related to inter-disciplinary collaborations are only starting to be explored systematically.

In my paper I focus on the problem of trust. According to Hardwig (1985, 1991), trust is essential for the epistemology of science; in fact, it is essential for *knowledge* in science. We can accept Hardwig's strong epistemic thesis, or recognize the fact that trust is certainly essential for the *success* of research (where success is a wider concept and can take a number of meanings: prediction, technological advancement, etc.). Either way, the problem I highlight in this paper is that trust networks are much weaker when it comes to interdisciplinary research teams. I explore the problem theoretically by considering two interpretations of trust: a Humean one, based on the concept of scientific evidence, and the one favored by Hardwig, based on ethical categories.

I then move on to exploring the problem on an empirical level. I use the worth of research conducted by MacLeod and Nersessian on two interdisciplinary research teams in Integrative Systems Biology (for short ISB, i.e. research groups where engineers collaborate with biologist to develop system explanations of biological phenomena as well as related technologies). Using cognitive ethnography (Hutchins 1995), MacLeod and Nersessian (2013a,b,c, 2014a,b) observed two ISB teams over a period of four years, collected materials, ran interviews and observed the interactions of interdisciplinary researchers while they collaborated over a number of research tasks. Their evidence confirms and strengthens the theoretical analysis of the problem of trust in interdisciplinary research.

The final section of the paper provides some tentative suggestions for strategies that can strengthen interdisciplinary trust. It suggests that philosophy, and philosophy of science in particular, has the tools needed for further exploration of these strategies.

Dupré, John. 2001. *Human Nature and the Limits of Science*. Oxford: Oxford University Press.

Hardwig, John. 1985. 'Epistemic dependence.' *The Journal of Philosophy* 82(7):335-349.

Hardwig, John. 1991. 'The role of trust in knowledge.' *The Journal of Philosophy* 88(12):693-708.

Kellert, Stephen H., Helen E. Longino, and C. Kenneth Waters. 2006. 'The Pluralist Stance.' In: Stephen H. Kellert, Helen E. Longino, and C. Kenneth Waters (ed.), *Minnesota Studies in the Philosophy of Science — Volume XIX: Scientific Pluralism*. Minneapolis: University of Minnesota Press: pp. vii-xxviii.

MacLeod, Miles, and Nancy J. Nersessian. 2013a. 'Building Simulations from the Ground Up: Modeling and Theory in Systems Biology.' *Philosophy of Science* 80(4):533-556.

MacLeod, Miles, and Nancy J. Nersessian. 2013b. 'The creative industry of integrative systems biology.' *Mind & Society* 12(1):35-48.

MacLeod, Miles, and Nancy J. Nersessian. 2013c. 'Coupling simulation and experiment: The bimodal strategy in integrative systems biology.' *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 44(4):572-584.

MacLeod, Miles and Nancy J. Nersessian. 2014a. 'Strategies for coordinating experimentation and modeling in integrative systems biology.' Forthcoming in *Journal of Experimental Zoology Part B*.

MacLeod, Miles and Nancy J. Nersessian. 2014b. 'Transdisciplinary Problem-Solving: Emerging Modes in Integrative Systems Biology.' Unpublished manuscript, source: Personal correspondence.

**Conor Mayo-Wilson. Structural chaos.** Philosophers of science often distinguish between two sources of predictive error: parameter error and structural model error (SME) (Parker 2010). One type of parameter error occurs when a researcher misidentifies the initial conditions of a dynamical system; call this initial conditions error (ICE). In a recent paper, Frigg et al. argue that the distinction between SME and ICE is crucial for both scientific practice and policy-making. They claim that, although there are methods that can generate accurate forecasts in the presence of both (i) ICE and (ii) chaos, there are no such methods for doing the same with respect to (i') SME and (ii') an analogous notion of 'structural chaos', which they call the 'hawk-moth' effect. For this reason, Frigg et al. argue that structural chaos and SME are neglected, but important topics within philosophy of science.

Although they provide an illustrative example and ample computer simulations to suggest structural chaos might be widespread, Frigg et al. neither define 'structural chaos' nor investigate the relationship between chaos (simpliciter) and structural chaos. This is important because there are dozens of definitions of 'chaos' in applied mathematics, and consequently, there might be dozens of analogous notions of structural chaos. At points, Frigg et al. seem to claim the structural chaos is the opposite of structural stability, which is intended to formalize the idea that small changes to a model do not result in large changes in its behavior. Unfortunately, there are several different definitions of 'structural stability' (Pugh 2008), and the relationships among the various notions of structural instability and different definitions of chaos are not known.

Frigg et al.'s research and related work on structural stability, therefore, raises at least three important questions for philosophers of science, applied mathematicians, and working scientists. First, for each definition of 'chaotic system', what is the analogous concept of structural chaos? Second, what are the relationships among the various notions of chaos (simpliciter), the analogous notions of structural chaos, and the various notions of structural stability? Finally, what are the implications of structural chaos for prediction, control, and explanation?

This paper takes a (very) preliminary step with respect to the first two questions. I first provide one candidate definition of 'structural chaos.' My definition is analogous to the notion of 'topological mixing', which is sometimes used to characterize chaotic systems. In particular, topologically mixing dynamical systems are chaotic according to a frequently-cited definition of chaos given by Devaney.

The central result of this paper is that, when a sufficiently rich class of models contains a chaotic function, then the entire class is structurally chaotic in my sense. The result is especially interesting because (1) a class of models may be structurally chaotic near the function  $F$ , in my sense, and yet (2)  $F$  may be structurally stable in several standard senses. That is, one obvious way of defining 'structural chaos' is logically consistent with standard concepts of structural stability. This raises the question of what the implications are of the various notions of structural stability for prediction, control, and explanation.

**James Nguyen. Why data models do not supply the target structure required by the structuralist account of scientific representation.** Scientific representation is now a firmly established topic in the philosophy of science. My preferred way of addressing the topic is with the following question: in virtue of what do scientific models represent their targets? The structuralist's answer to this question is that scientific models are mathematical structures that represent their targets, in least in part, in virtue of a certain morphism between the two. All variations on the structuralist account of scientific representation share one fundamental assumption: there is a target-end structure to be represented. This assumption is defended in two different ways: target systems instantiate structures, or alternatively, it is data models that supply the target-end structures required. In this paper I investigate the latter approach.

I identify three objections that any data model structuralist must face: a naturalistic objection that appeals to the face-value content of actual scientific models, the further objection that this content really is about physical systems and not data, and finally the loss of reality objection, that data model structuralists miss the world. Van Fraassen's *Scientific Representation: Paradoxes of Perspective* provides an intriguing argument in response to at least one of these objections. Unfortunately it has yet to receive the attention it deserves. I suspect this is because of the considerable novelty, at least in the context of scientific representation, of the exciting notions appealed to: indexicality, location in logical space, representation as thus-and-so, and pragmatic tautologies amongst others. And the fact that the argument is spread across Part I–III of the book, interwoven with substantial broader discussions of representation, measurement, and empiricism.

In this paper I reconstruct van Fraassen's argument for the claim that, in a given context, for an individual scientist, there is no difference between representing a target system and representing data gathered from the system. I argue that, if successful, the claim would succeed in meeting at least one of the objections to data model structuralism, and possibly others as well. Unfortunately, at least for the structuralist, I argue that van Fraassen's argument fails at two crucial junctures. Firstly, he equivocates between location and self-location in logical space. And secondly, he equivocates between representing a target as thus-and-so and asserting that a target is thus-and-so. As such I conclude that van Fraassen's response to the objections raised fails. The question remains: where are the target-end structures, necessary for any structuralist account of scientific representation, to be found?

**Robert Northcott & Anna Alexandrova. Armchair science.** There seems to be something especially wrong with the idea of armchair science, more so than armchair philosophy. Armchair science – which, we argue, is what a large part of modeling in the social and biological sciences amounts to – endeavors to find truths about the world without observational or experimental input. We define it as modeling whose goal is neither prediction nor more or less faithful representation. Instead it exists, as it were, at one remove from direct empirical contact, exploring relations of dependence between variables that may bear only a loose relation to real-world entities.

Much of theoretical economics and political science – especially the rational choice project – and parts of theoretical biology are prime examples of armchair science. The contrast is with models in engineering, econometrics, climate science, election prediction and many other fields that are targeted at particular

real-world systems, and that are judged by successful prediction or explanation or other empirical criteria, rather than exclusively by theoretical innovation.

Although there is no sharp line between the two types of modeling, the category of armchair science is useful because it allows us to raise what we call the efficiency question. Of course, spending time in the armchair might be useful sometimes. But how useful? And compared to what? We make two claims:

1. There is a prima facie case that there is too much armchair science.
2. Philosophers of science have given armchair science too easy a ride.

We support the first claim with a case study of the modeling industry around the Prisoners Dilemma game. This literature, spanning economics, biology and even neuroscience, is large and impressive in the theoretical innovations it offers. However, given the resources devoted to it, its empirical pay-off is disappointing. True, it is claimed to apply to many phenomena, but we argue that this appearance is deceptive once one looks at the details.

We support the second claim by noting that philosophers of science have devoted much attention to armchair models, and by analyzing the implicit defenses of them that have thus emerged. The most common of these are: (1) armchair science is the best we can do when we cannot experiment, for it allows us to explore causal relations in the abstract; (2) idealized models can be de-idealized by robustness analysis; (3) idealized models can provide understanding even in the absence of predictive or explanatory success. All of these defenses, we argue, turn out to be dubious. But in addition to that, in any case none of them directly addresses the efficiency question because none addresses the central issue of opportunity cost – what could scientists be doing instead of, or at least in addition to, such armchair modeling?

We conclude by proposing how the efficiency question could be tackled. Since the main contrasts to armchair modeling are field and experimental methods, it requires judging the counterfactual of what pay-off would be gained if these alternative methods were adopted.

**Rune Nyrup. Analogical reasoning and pursuitworthiness.** John Norton has proposed a material account of analogical reasoning in science according to which there are no general inference rules or universal principles to justify analogical inferences. Rather, these can only be validated by local facts about analogies, i.e. systematic similarities, holding between the domains investigated.

While this account is plausible in cases where scientists already know whether the similarities (and differences) relevant to the inference obtain, it faces a problem in cases where such knowledge is lacking. Here, according to Norton, scientists can conjecture that an analogy obtains, the accuracy and scope of which can then be investigated empirically. The problem, which was also raised by the ‘Campbellian’ in Mary Hesse’s dialogue in *Models and Analogies in Science*, concerns why an analogy, rather than any other hypothesis, should be conjectured in this case. As it stands, the material account fails to account for the role analogical reasoning seems to play in scientific deliberation about which hypotheses to investigate. In terms of the distinction drawn by Larry Laudan, Norton’s account only captures the role of analogy in reasoning about which hypotheses to accept but not reasoning about which hypotheses to pursue.

Campbell himself insisted, in *Foundations of Science*, that theories would be completely ‘valueless and unworthy of the name’ (p. 129) without analogies, since these distinguish theories ‘from the multitude of others [...] which might also be proposed to explain the same laws’ (p. 142). Paul Bartha has used this remark as inspiration for an account of analogical reasoning according to which analogies can show a hypothesis to be prima facie plausible. In Bartha’s interpretation, this involves providing reasons to think the hypothesis might be true. This interpretation, I argue, is mistaken. Firstly, Bartha’s formal arguments are too weak to establish that analogies can, in general, provide even this type of support. Insofar as analogical reasoning show hypotheses any more likely, this is better captured by Norton’s material account.

Secondly, though a hypothesis needs to be minimally plausible to be worth pursuing, this is not sufficient. I propose an account according to which justification for pursuing a hypothesis H stems from weighing the potential epistemic gains of learning whether H is true, and the likelihood that pursuing H will allow us to learn this, against the expected costs of doing so. Thus, showing H worthy of pursuit can stem from reasoning showing that it would more valuable to learn whether H is true than previously thought, rather than from showing it more likely.

I argue that analogical reasoning can work in exactly this way, thus interpreting more literally Campbell’s claim that analogies pertain to the value of theories. Firstly, analogous hypotheses tend to provide a high degree of explanation and understanding, both widely recognised as important epistemic values. Secondly, my account is consistent with Bartha’s pragmatic argument that analogical reasoning provides a good balance between conservative and innovative epistemic values, but without having to presuppose any further connection to the truth or likelihood of the hypothesis.

**Matthew Parker. The poverty of infinitesimal probabilities.** Several philosophers have argued that probabilities (chances or rational credences, depending on the author) should be regular, i.e., only impossible events should have probability zero. This is the chief motivation for introducing infinitesimal and hyperreal probabilities, but such probabilities do not overcome all of the problems with regularity. As a few authors have noted, regular probabilities cannot be translation invariant. In fact, if point sets in the real line are assigned regular probabilities, then certain simple, countable, bounded, disjoint sets that are translations of each other must differ in probability. Few have regarded such translation variance as a serious problem for regularity, but it is. In the first place, it means that regular chances are not determined by space-time invariant laws and circumstances. So two outcomes of an experiment must differ in probability for no physical reason at all. Some have argued that the difference can be accounted for by ‘non-local’ factors, such as one set of outcomes being a proper subset of the other, but in our examples, the outcome sets are bounded and disjoint. Secondly, it means that, however symmetric our knowledge and evidence, regular credences cannot be symmetric. Lacking any evidence whatsoever to favour one outcome over another, it is rational to assign them equal credence. But regularity implies that, in some such cases, we must nonetheless give one outcome higher credence than the other. Some have pointed out that regular probabilities can be very nearly translation invariant, up to an infinitesimal difference, but this is no help if we want probabilities to represent fully accurate chances or fully symmetric evidence. On the other



hand, if we are content with very nearly representative probabilities, we have no need for infinitesimals in the first place.

**Makmiller Pedroso. The evolution of transient individuals.** This paper is concerned with the evolution of individuality—i.e., the evolution of stable collectives (e.g., multicellular organisms) from formerly independent units (e.g., single cells). Individuality is often understood as a social phenomenon. Stable individuals evolve because their parts cooperate and the chance for internal conflict is constrained. Prominent accounts of biological individuality propose two mechanisms to account for the evolution of stable individuals: reproductive bottlenecks and division of labor. I use bacterial communities called ‘biofilms’ to show the existence of ecological analogs for these two mechanisms: ecological bottlenecks and ecological specialization. Like their non-ecological counterparts, these ecological mechanisms can account for the stability of individuals because they increase the costs of cheating among the parts of a collective. A biofilm undergoes an ecological bottleneck when its population size drastically decreases because of mass-mortality events caused by, for example, antimicrobial treatments. Ecological bottlenecks decrease the frequency of cheats in a biofilm by increasing the genetic relatedness among its cells. Ecological specialists are produced through disruptive selection when biofilms grow in heterogeneous environments. At sufficiently high levels of niche specialization, the resistance to the occurrence of cheats increases within a biofilm. Both ecological bottlenecks and niche specialization are contingent upon extrinsic factors. Ecological bottlenecks depend on mass-mortality events; ecological specialization requires heterogeneous environments. These mechanisms suggest that individuality can be a transient phenomenon: a collective can be an individual in certain environments but not in others. This is a welcome consequence given that the evolution of individuality is a gradual process.

**Zee Perry. Intensive and extensive quantities.** Quantities are properties and relations which exhibit ‘quantitative structure’. For physical quantities, this structure can impact the non-quantitative world in different ways. In this paper I introduce and motivate a novel distinction between quantities based on the way their quantitative structure constrains the possible mereological structure of their instances. I borrow the terms ‘extensive’ and ‘intensive’ for these categories, though my use is substantially revisionary. I present and motivate this distinction using two case studies of successful physical measurements. (of mass and length, respectively). I argue that the best explanation for the success of the length measurement requires us to adopt my notion of extensiveness. I consider and reject an alternative to extensiveness, commonly called ‘additivity’. I demonstrate that this new distinction can do what the additive/non-additive distinction cannot in explaining the success of length measurements. I also briefly sketch further applications of the intensive/extensive distinction, specifically to the project of giving a satisfactory account of quantitative structure in non-mathematical and non-metrical terms.

**J. Brian Pitts. Real change in Hamiltonian General Relativity.** The Earman–Maudlin standoff over change in Hamiltonian General Relativity calls for re-examination of the Hamiltonian formalism. This work continues the recent trend from Mukunda, Castellani, Sugano, Pons, Salisbury, Shepley

and Sundermeyer toward recovering Lagrangian equivalence in Dirac-Bergmann constrained Hamiltonian dynamics, which was present in the earliest works but later lost. A first-class constraint typically does not alone generate a gauge transformation, contrary to widespread belief: by direct calculation it is found that each first-class constraint in Maxwell’s theory generates a change in the electric field  $E$  by an arbitrary gradient, spoiling Gauss’s law. The secondary first-class constraint  $p^i{}_{,i} = 0$  still holds, but being a function of derivatives of momenta, it is not directly about  $E$  (a function of derivatives of  $A$ ). The canonical momenta  $p^i$ , being auxiliary fields, acquire physical meaning only parasitically on the velocities using Hamilton’s equations  $\dot{q} - \frac{\partial H}{\partial p} = -E_i - p^i = 0$ . Only a special combination of the two first-class constraints, the Anderson-Bergmann (1951)-Castellani gauge generator  $G$ , leaves  $E$  unchanged, preserves Hamilton’s equations, and plays the expected role in Noether’s second theorem.

An error in Dirac’s argument that a primary first-class constraint generates a gauge transformation is noted: he cancels out any possible effects on initial data by comparing evolution from two identical configurations. Hence the Dirac conjecture that secondary first-class constraints also generate gauge transformations, cannot even get started due to a false presupposition. The usual concept of Dirac observables should also be modified to employ the gauge generator  $G$ , not the first-class constraints separately, so that the Hamiltonian electromagnetic observables become equivalent to the Lagrangian ones such as  $F_{\mu\nu}$ .

In General Relativity in Hamiltonian form, change has seemed to be missing, defined only asymptotically, or otherwise obscured at best, because the Hamiltonian is a sum of first-class constraints and a boundary term and thus supposedly generates gauge transformations. Once one knows that only the gauge generator  $G$ , not an individual first-class constraint, generates a gauge transformation, the problem of missing change disappears. Examining Hamilton’s equations in a toy theory discarding spatial variation from GR for simplicity, one finds that there is time dependence in the Hamiltonian formalism in all time coordinatizations for solutions if and only if there is no time-like Killing vector field: the Hamiltonian and 4-dimensional differential geometric criteria for change agree. The inclusion of a massive scalar field is simple. No obstruction is expected in including spatial dependence and coupling more general matter fields. Hence change is real and local even in the Hamiltonian formalism. Distinguishing internal and external symmetries leads to a revised Lagrangian-equivalent definition of observables in GR as just geometric objects, hence varying over time and space. The considerations here resolve the Earman–Maudlin standoff: the reformed Hamiltonian formalism does not have absurd consequences for change and observables. Hence the classical part of the problem of time in canonical quantum gravity is resolved. A further issue involving quantum constraints is not addressed.

**Dave Race. Filling in surplus structure in the partial structures framework.** A supposed advantage of the partial structures framework is that it makes the notion of surplus structure more precise. Physical theories are embedded into mathematical structures, which grants the theory access to ‘surplus’ mathematical structure. By understanding surplus structure in terms of a family of structures, the framework is used to represent the links between the structures, and between the mathematics and the physics (French, 187-207, CUP, 1999). The framework’s ability to perform this role is said to be due to the R3-component of the partial structures capturing the ‘open-



ness' of scientific theories. This advantage has been overstated due to inconsistent claims over the role of the R3-component, vagueness over what it is supposed to contain, and equivocation over surplus structure. These problems can be mostly resolved by distinguishing four types of surplus structure and locating these in different areas of the partial structures version of the semantic view (SV).

Redhead (Synthese 32: 77–112, 1975; 73–90, CUP, 2001) outlines four distinct types of surplus structure: uninterpretable mathematics; uninterpreted mathematics, subsequently interpreted; certain idealisations; and the accommodation of Hesse's account of analogies. The SV uses partial structures to accommodate 'vertical' relationships between theoretical structures and data models', idealisations, 'horizontal' inter-relationships between theories', relationships involved in theory change and construction (French, 2000), and in the Inferential Conception of mathematical applicability (Bueno & Colyvan, *Noûs* 45(2):345-374, 2011; Bueno & French, *BJPS* 63: 85–113, 2012), which holds mathematics is applicable because of partial morphisms holding between empirical set ups and mathematical models. One maps from an empirical set up to the mathematical domain via an immersion mapping; one then draws consequences from the mathematical formalism; the resultant structure is then interpreted in terms of the initial empirical set up via an interpretation mapping. The immersion mapping can be repeated, allowing for the embedding of a mathematical structure in another.

I will argue that the analogy type of surplus structure and the 'mismatch' role for the R3-component should be isolated to non-IC parts of the SV, and that the second and third types occur during the interpretation mapping of the IC. I claim that the SV is committed to an 'as if' interpretation of idealisations, and that the first type of surplus structure involves uninterpreted families of mathematical structures that provide new inferential relations and should occur during an immersion mapping from one mathematical structure to another. The SV faces a serious problem in reconciling these commitments, which can be introduced as the answer to the question 'how is surplus structure to be introduced?' I argue that in attempting to answer this, the proponents of the SV face a dilemma: either the partial structures in the derivation step must contain all mathematical structures (in the R3-component); or that the mathematics contained in the R3-component must be restricted in some justified way, but that no such justification is available. I sketch some possible solutions to this dilemma, and conclude that the dilemma might be avoidable if some of these options are developed.

**Grant Ramsey & Charles Pence. Is organismic fitness at the basis of evolutionary theory?** Fitness is a central theoretical concept in evolutionary theory. Despite its importance, much debate has occurred over how to conceptualize and formalize fitness. One point of debate concerns the roles of organismic and trait fitness. In a recent addition to this debate, Elliott Sober (2013) argues that trait fitness is the central fitness concept in evolutionary biology, and that organismic fitness is of little value.

Sober's argument has much to recommend it. First and foremost, his clarity regarding the distinction between individual fitness and the fitness of traits, as well as the relationship between the two, has been sadly lacking in recent literature on fitness. But we will argue here that his central thesis – that individual fitness is broadly irrelevant – is mistaken, and that this mistake arises as a result of confusion over the variety of roles that the notion

of fitness plays in evolutionary theory. While trait fitness is the salient concept for some of the roles of fitness, it is individual fitness that is foundational.

Many of the most important uses of fitness fall under two categories. First is what we will call a metrological role of fitness – that is, fitness's role as a quantitative measure in evolutionary studies. Biologists can measure the realized fitness of organisms by tallying such things as their lifetime reproductive success, and they can measure trait fitness by recording trait changes over time. Second is what we will call the conceptual role of fitness – that is, fitness as an element of the causal or explanatory structure of evolutionary theory.

Keeping this distinction in mind, then, our argument proceeds as follows. We begin by arguing that there exist three common conceptions of trait fitness – and each of these, in turn, is parasitic on individual fitness, making individual fitness the fundamental notion of fitness in the conceptual role. Next we argue that in the metrological role, the situation is less clear – there are certainly studies in which trait fitness is the more important concept. But it is, we claim, far from true that, as Sober argues, 'evolutionary biology has little use for [individual] fitness' (2013, p. 336). In a wide variety of examples, we argue, it is indeed the fitness of individual organisms that biologists look to measure, even when they make inferences about the fitness of traits from those measurements. Individual fitness is therefore fundamental in the conceptual role, and useful in the metrological role, and should thus, contra Sober, by no means be rejected outright.

Sober, Elliott. 2013. 'Trait Fitness Is Not a Propensity, but Fitness Variation Is.' *Studies in History and Philosophy of Biological and Biomedical Sciences*. 44: 336–41. doi:10.1016/j.shpsc.2013.03.002.

**Alexander Reutlinger. What's explanatory about non-causal explanations?** According to the causal model of explanation, the sciences explain by providing information about causes and causal mechanisms (cf. Salmon 1984, Machamer, Darden and Craver 2000, Woodward and Hitchcock 2003, Strevens 2008). Causal models are among the most widely accepted models of explanation today. However, in the past decade, an increasing number of philosophers have argued that the explanatory practices in the sciences are richer than causal model of explanation suggests (cf. Batterman 2002, 2010, Lipton 2004, Lange 2010, 2012, Pincock 2012). These philosophers claim that there are non-causal explanations that cannot be accommodated by the causal model. Case studies of non-causal explanations come in a surprisingly diverse variety: for instance, the non-causal character of scientific explanations is based on the explanatory use of non-causal laws, purely mathematical facts, symmetry principles, inter-theoretic relations, renormalization group methods, and so forth. If there are instances of non-causal ways of explaining, then the causal model, at least, cannot be the whole story about scientific explanation. However, the natural follow-up question why non-causal explanations are explanatory is seldom addressed. That is, it is seldom asked which (if any) philosophical model of explanation adequately describes non-causal explanations.

My goal in this talk is to provide an answer to this question: that is, I will advocate a philosophical model of non-causal explanations in the sciences. My main claim is that non-causal explanations can be understood by extending received causal difference-making accounts of scientific explanations to a generalized difference-making account: non-causal explanations re-

veal non-causal counterfactual dependencies between explanandum and explanans, or so I will argue.

More precisely, I will argue that a generalized model can be obtained by amending Woodward and Hitchcock's (2003) notion of explanatory difference-making, which – in essence – boils down to counterfactual dependencies between the explanandum and the explanans. These counterfactual dependencies are revealed in answers to so-called 'what if things had been different' questions. I argue that counterfactual dependencies need not be given a causal interpretation. Rather such dependencies cover a broader range than merely causal dependencies, as Woodward suggests in a side remark but does not elaborate: 'the common element in many forms of explanation, both causal and non-causal, is that they must answer what-if-things-had-been-different questions.' (Woodward 2003: 221).

I will elaborate this idea that non-causal explanations work by providing information about non-causal counterfactual dependence relations (here, I am building on the work of Bokulich 2008, Reutlinger 2011, 2013, Saatsi and Pexton 2013, Pincock forthcoming). It is my aim to show that, at least, certain kinds of non-causal explanation can be explicated as answers to 'what if things had been different questions'. My main case studies supporting this claim are (a) applied mathematical explanations, (b) renormalization group explanations of universality, and (c) explanations relying on symmetry principles. Finally, I will conclude by arguing that a difference-making account of non-causal explanations is a serious and preferable alternative to Marc Lange's (2012, 2013) 'modal' account, according to which non-causal explanations work by showing that the explanandum had to occur.

**John Roberts. Humean laws and explanation.** One standard objection to 'Humean' theories of laws is that such theories make it impossible to account for the power of laws to explain their instances. For Humeans, the lawhood of the laws is constituted by patterns in the great mosaic of local, non-modal states of affairs, so they are ultimately constituted by those states of affairs; how, then, can they explain one of those states of affairs without circularity? Loewer has argued that this objection fails because there is an important distinction between metaphysical explanation and scientific explanation, and the Humean mosaic metaphysically explains the laws whereas the laws scientifically explain particular matters of fact, so there is no vicious explanatory circle. Lange has recently argued that Loewer's reply is unsuccessful. I agree with Lange.

Here, I consider a number of other possible replies to the objection. For example:

(1) From the Humean point of view, covering-law explanation is importantly analogous to certain sorts of aesthetic explanation, wherein the role of a certain element in a work of art is explained by the arrangement of all the work's elements (e.g. one action of a character in a novel is made sense of by reference to a larger pattern in the novel, of which that action is partly constitutive; the occurrence of a phrase at a certain point in a movement is explained in part by reference to the fact that the movement is of sonata form, a fact of which the occurrence of that phrase at that point is partly constitutive). One might reply that in such aesthetic explanations, what is really doing the explaining is not the larger pattern of which the element is partly constitutive, but rather the intention of the artist to create a work exhibiting that pattern, and this wrecks the analogy with covering-law explanations. In reply, I agree that one could understand such aesthetic explanations in this way, but this is not the only way to under-

stand them. On another way, such aesthetic explanations are not causal explanations – similarly, a Humean might argue that despite superficial appearances to the contrary, covering-law explanations are not best understood as causal explanations.

(2) The law in a covering-law explanation is perhaps not properly speaking part of the explanans, but rather is a principle that accounts for why the initial conditions cited are capable of explaining the explanandum (e.g. the earth's orbit is explained by the sun's gravitational force on it, and Newton's second law of motion merely accounts for why the latter is capable of explaining the former).

(3) Compare: Modus ponens plays a role in many scientific explanations; what constitutes its truth? Arguably, lack of counterexamples to MP throughout the set of actual facts – which includes the explainanda of many explanations in which MP plays a role.

(4) Perhaps the laws supervene on the Humean base, but (as argued in [reference deleted]) their lawhood (and so, their explanatory power) derives from their normative status, which is not reducible to the Humean base.

**Carlo Rossi. Enduring a relativistic world.** Endurance theorists have often appealed to the notions of exact location or occupation and multi-location in order to explain how objects persist through spacetime in the context of the Special Theory of Relativity (STR). Specifically, endurantists invoke these two notions in order to claim that objects persist through spacetime by exactly occupying multiple spacetime regions, each of which is temporally unextended and disjoint from the other. The aim of this paper is to provide a better understanding of these two notions and of the implications they have for our preferred account of endurance. Bearing such aim in mind, in the first section of the paper I discuss the five conditions proposed by Cody Gilmore that any account of exact occupation must satisfy, and also the difficulties that arise for this cluster of conditions (2006). In the next section I evaluate Parsons' alternative proposal, which defines exact occupation in terms of overlap (2007). In spite of some advantages over Gilmore's account, one noticeable shortcoming of this account is that it does not allow enduring objects to be multi-located at different spacetime regions. Enduring objects exactly occupy one spacetime region, which coincides with their spatiotemporal path. Next, I explore the possibility of a middle ground between Gilmore's and Parsons' account, which might allow us to retain the advantages of Parsons' accounts along with multi-location. Such theory seems to be defended by Crisp and Smith (2005), but I argue that they fail in their attempt of treating overlap as primitive and at the same time allowing multi-location. If time allows, I will finally discuss the prospects for some alternative ways of characterizing the endurance vs. perdurance debate which are available for those who remain skeptics of the intelligibility of the notion of multi-location. Crucially, these ways of characterizing the current debate would switch its focus of the dispute from issues about location to issues about parthood (Donnelly 2010, 2011.)

**Juha Saatsi. Worthwhile distinctions: Kinematic, dynamic, and (non-)causal explanations.** It is commonplace nowadays to accept that some scientific explanations are non-causal. Very little has been said in general terms of such non-causal explanations, however, and some philosophers still stubbornly defend the hegemony of causal explanations. (e.g. Skow, B. 'Are There Non-Causal Explanations?', BJPS 2013.)

On the whole there is really no agreement as to how to demarcate between causal and non-causal explanations.

In this paper I examine and throw light on this issue of causal/non-causal demarcation from the perspective of a related (but not identical) distinction in physics: kinematic vs. dynamical explanations. A secondary objective is to respond to those who (with Skow) maintain that all explanations (of particular facts) are causal. I argue that defending the hegemony of causal explanation in the philosophy of science risks failing to recognize distinctions that are worthwhile not only to philosophers of science, but to scientists themselves.

The fact that physicists themselves draw and habitually employ a distinction between kinematic and dynamic explanations is a clear indication of it being a substantial, worthwhile distinction to draw. A typical (but by no means definitive) textbook presentation of this distinction refers to:

(1) Kinematics as the study of the geometry of motion; kinematics is used to relate displacement, velocity, acceleration, and time, without reference to the cause of the motion.

(2) Dynamics as the study of the relation existing between the forces acting on a body, the mass of the body, and the motion of the body; dynamics is used to predict the motion caused by given forces or to determine the forces required to produce a given motion.

(cf. Beer, F. *Vector mechanics for engineers: statics and dynamics*, 2009.)

The above distinction between kinematics and dynamics suggests that explanations belonging to kinematics are non-causal, while those belonging to dynamics are causal. Delineating the contrast between kinematical vs. dynamical explanations is far from straightforward, however. I will start with a brief survey of the kinematic vs. dynamic distinction as employed in physics, starting from its historical origins, and noting some of ambiguities and variances in its use. I will then proceed by providing exemplars of kinematic explanations.

The first exemplar involves a familiar, simple and intuitive powerful toy-model of a paradigmatic kinematic explanation in order to explore the sense in which a kinematic explanation can furnish a genuinely non-causal explanation by virtue of attending to the 'geometry of motion' (cf. (1) above).

The second exemplar involves a real scientific model from quantum mechanics that furnishes a non-causal, kinematic explanation of the behaviour of a fermionic many-particle system. The involves the Pauli exclusion principle in a way that is particularly apposite and interesting in the context of the philosophical literature, which has persistently made wrongheaded claims about the relevance of the exclusion principle in relation to the causal/non-causal debate (e.g. Skow, *ibid.*).

**Ryan Samaroo. There is no conspiracy of inertia.** I examine two claims that arise in Harvey Brown's account of inertial motion in *Physical Relativity* (2005). Brown claims there is something objectionable about the way in which the motions of free particles in Newtonian theory and special relativity are coordinated. Brown also claims that since a geodesic principle can be derived in Einsteinian gravitation the objectionable feature is explained away. I argue that there is nothing objectionable about inertia and that, while the theorems that motivate Brown's second claim can be said to figure in a deductive-nomological explanation, their main contribution lies in their explication rather than their explanation of inertial motion.

I begin by examining Brown's claim that there is something objectionable—something conspiratorial—about inertia

in Newtonian theory and special relativity. I argue that the alleged conspiracy is motivated by a commitment to a number of metaphysical principles or intuitions that have their source in Einstein, and that the allegation obscures rather than clarifies inertia in Newtonian theory and special relativity. I argue that all of these intuitions are bound up with an underlying view according to which there is something objectionable about absolute or global space-time structures.

I also consider an implication of accepting Brown's claim that there is something conspiratorial about inertia: I consider the suggestion that, for a conspiracy theorist, all physical theories are conspiratorial. But I argue that, if it is a view about absolute or global space-time structures that is driving the alleged conspiracy, then there is little to gain by explaining away the conspiracy of inertia by appealing to Einsteinian gravitation, for one can point to conspiratorial features even in that framework.

In the second part of the talk, I address Brown's claim that inertia is explained by Einsteinian gravitation because a geodesic principle can be derived from the field equations. I review Weatherall's (2011) challenge to Brown's claim. Weatherall argues that, if there is any sense in which Einsteinian gravitation can be said to explain inertia, then geometrised Newtonian gravitation explains it at least as well. While I agree with Weatherall, I argue that there is a better way of thinking about the geodesic theorems. Their main contribution lies not in their *explanation* of inertial motion but in their *explication* of it. I argue that the geodesic theorems of Geroch and Jang (1975) and Weatherall (2011) explicate inertial motion by making perspicuous the dependency of inertial motion on the conservation of momentum. This is manifest, though under-appreciated, in Newton's own account of inertia, and I argue that the work of his successors—notably, d'Alembert, Thompson and Tait, and Maxwell—represents a deliberate attempt to establish the fundamental importance of the conservation principle. In spite of their important differences, old-fashioned Newtonian theory, geometrised Newtonian gravitation, and Einsteinian gravitation are strongly analogous in their accounts of inertial motion.

**Mario Santos-Sousa. What, if anything, can the epistemology of number learn from the psychology of numerical cognition?** My goal is to address the question raised in the title of this paper: What, if anything, can the epistemology of number learn from the psychology of numerical cognition? I plan to proceed as follows. First, I will ask whether the psychology of numerical cognition has anything to contribute to the epistemology of number, which I shall answer in the affirmative. Next, I will identify its specific contribution in the light of contemporary research on numerical cognition.

I begin by considering the following problem, which is distinctively epistemological: How is it possible for us to know about numbers? This problem gains its bite from conflicting assumptions about the knowledge in question. One is a set of specific assumptions about its subject matter: (M1) numbers are mind-independent entities, and (M2) numbers are abstract entities. The other is a set of general assumptions about its possible sources: (E1) experience is the ultimate source of our knowledge of mind-independent entities, and (E2) experience cannot yield knowledge of abstract entities. It is not difficult to see how these assumptions, taken together, would make it impossible for us to know anything about numbers. Herein lies the force of the challenge.

The philosophical literature abounds with attempts at meeting this challenge, either by rejecting M1 or M2 (or both),

or by rejecting E1. Here, I will focus on E2 instead, which has received comparatively less attention. In particular, I will defend that experience can sometimes yield knowledge of abstract entities such as numbers. In order to see how, we will have to look into the psychological literature on numerical cognition. I will focus specifically on our knowledge of cardinal numbers, which, for present purposes, I shall take to be properties of sets or collections (if we allow for collectives of one and zero items).

The available evidence suggests that we have an innate capacity for detecting the cardinal size of collections of perceptually presented items. What's crucial to my present argument is that this number sense does not seem to be tied to any specific sensory modality, such as vision, but tracks numerical information reliably across different modalities (auditory, tactual, etc.). In other words, the capacity in question allows us to experience cardinality 'as such'.

However, this capacity is very limited. It only comprises a rough sense of large cardinal size and an exact sense of small cardinal size: we discriminate small collections of objects (but only up to a certain threshold of about four) and are able to approximate larger numerical quantities (but fail to capture numerical differences below a certain ratio). So, how do we transcend these limitations? The obvious answer is: we count. Counting is an experiential means of obtaining information about cardinal numbers that we wouldn't otherwise be able to discriminate. Moreover, the standard counting principles yield information about cardinality 'as such' irrespective of the individual natures of the items being counted.

I therefore conclude that we can gain knowledge of numerical abstracta through our number sense and counting experience.

**Arianne Shahvisi. Eliminating conspiracies via the genealogy of subsystems.** Fodor (1997) suggests that elaborate conspiracies amongst fundamental particles are required in order to bring about the projectibility of special science generalisations (SSGs). Callender and Cohen (2010) undertake to debunk this conspiracy, arguing that (a) Albert (e.g. 2000) and Loewer's (e.g. 2008) theory—which sees SSGs as probabilistic corollaries of the fundamental laws plus the Past Hypothesis and Statistical Postulate—does not succeed in accounting for the projectibility of SSGs, and that (b) their own relativised Mill–Ramsey–Lewis theory of lawhood—the 'Better Best System' (BBS)—is the most effective available solution. In this paper I challenge them on both points, arguing that a synthesis of aspects of theirs and Albert and Loewer's theories is necessary in order to decisively rule out the conspiracy and simultaneously respect the autonomy of the special sciences. Specifically, I will find that the optimal non-conspiratorial theory of lawhood is one that considers the way in which the origins of macroscopic subsystems restricts their later behaviour. I call this the 'Subsystem Genealogy' amendment, and propose that it could close vital explanatory lacunae in the otherwise powerful BBS theory, or could alternatively be seen as a conceptual version of the AL theory, which bypasses the probabilistic details. By considering the origins of these macroscopic subsystems, BBS can transcend its status as a merely descriptive account of special science lawhood, and set its sights on also explaining the projectibility of SSGs. Since BBS has considerable merits as a theory of lawhood—notably the ability to combat challenges to counterfactual asymmetry, such as Elga's (2000)—it is especially important that the mechanism, as well as the content, of the theory is understood.

Albert, D. Z., (2000), *Time and Chance*, Harvard University Press.

Callender, C. & Cohen, J. (2010) 'Special Sciences, Conspiracy and the Better Best System Account of Laws', *Erkenntnis*, 73: 427-447.

Elga, A. (2000). 'Statistical Mechanics and the Asymmetry of Counterfactual Dependence.' *Philosophy of Science suppl.* Vol 68: 313-24.

Fodor, J. (1998). 'Special Sciences; still autonomous after all these years', *Philosophical Perspectives*, 11, 149-163.

Loewer, B. (2008). *Why There is anything except physics. In Being reduced: new essays on reduction, explanation, and causation* (eds H. Jakob & K. Jesper). New York, NY: Oxford University Press.

**Mauricio Suárez. Probabilistic dispositions, chance distributions, and experimental statistics.** Probabilistic modelling may be most generally described as the attempt to characterise (finite) experimental data in terms of models formally involving probabilities. I argue that a coherent understanding of much of the practice of probabilistic modelling calls for a distinction between three notions that are often conflated in the philosophy of probability literature. A probability model is often implicitly or explicitly embedded in a theoretical framework that provides explanatory – not merely descriptive – strategies and heuristics. Such frameworks often appeal to genuine properties of objects, systems or configurations, with putatively some explanatory function. The literature provides examples of formally precise rules for introducing such properties at the individual or token level in the description of statistically relevant populations (Dawid 2007, and forthcoming). Thus, I claim, it becomes useful to distinguish probabilistic dispositions (or single-case propensities), chance distributions (or probabilities), and experimental statistics (or frequencies). I illustrate the distinction with some elementary examples of games of chance, and go on to claim that it is readily applicable to more complex probabilistic phenomena, notably quantum phenomena.

I then argue that it is possible to understand the role of these three notions in probabilistic modelling in terms of Bogen and Woodward's (1988) three-tier or tripartite distinction between theory, phenomena and data. Thus I suggest that in the context of probabilistic modelling, propensities are best understood as explanatory posits of theory, which both ground and explain chance or probability distributions. These distributions in turn are often best understood as models of phenomena in the sense described by Woodward and Bogen. Finally, relative frequencies of particular experimental outcomes in a given sequence constitute experimental data. It follows from the application of this tripartite distinction that propensities are typically not to explain particular outcomes or experimental data but rather the phenomena in the form of chance or probability distributions. The statistical data in turn may be used to directly confirm (and therefore also to test) probabilities, but not propensity ascriptions. The ascription of particular propensities – as Charles Peirce noted long ago, see also Mellor 2013 – is rather to be justified (or criticized) by abductive means in terms of their explanatory qualities.

I finally briefly review arguments in favour of similar conceptual distinctions within the philosophy of probability literature. I find that there are good philosophical reasons – independent of the considerations from modelling practice reviewed above, and related instead to arguments for objective chance

found in Mellor, 2005 – that already recommend the distinction between propensities, probabilities and frequencies.

Bogen, J. and J. Woodward (1988), 'Saving the Phenomena', *The Philosophical Review* XCVII (3), 303-352.

Dawid, A. P. (2007), Counterfactuals, hypotheticals and potential responses: A philosophical explanation of statistical causality, in F. Russo and J. Williamson (eds.), *Causality and Probability in the Sciences*, London College Texts, pp. 503-532.

Mellor, H. (2005), *Probability: A Philosophical Introduction*, London: Routledge.

\_\_\_\_\_ (2013), 'Propensities and Pragmatism', *The Journal of Philosophy*, vol. CX (2), 61-92.

**Piotr Szalek. The Duhem–Quine Thesis reconsidered.** The high point of falsification of physical theories in a standard view of philosophy of science is the so-called crucial experiment. It is a kind of manipulated empirical test, which provides the criterion for distinguishing between two rival hypotheses, where one is an acceptable theory due to a pass of the test, and the other turns out to be an unacceptable theory as it does not pass the test.

The crucial experiment was supposed to play its significant role, because in virtue of an empirical disconfirmation of one theory, the experiment was assumed to confirm the other as true. However, in the 'La théorie physique, son objet et la structure' (hereafter quoted in English translation as 'The Aim and Structure of Physical Theory' ([1906/1954]), Pierre Duhem famously argued against this view and stated that crucial experiments in physics are impossible as they are necessary ambiguous and logically incomplete. His contention rested on the claim that '[a] physical theory is not an explanation [of the true reality in itself in virtue of some broad metaphysical ramification of physics]. It is a system of mathematical propositions, deduced from a small number of principles, which aim to represent as simply, as completely, and as accurately as possible a set of experimental laws' (Duhem [1906/1954], 19). Furthermore, the different theories could be equally suitable to represent a given group of experimental laws. And, assuming holism, no hypothesis could be tested in isolation, but merely as a part of an entire set of scientific theory. The problem which Duhem identified in 1906 was slightly overshadowed and neglected in a mainstream philosophy of science until a challenging paper by Willard Van Orman Quine published in 1953 and entitled 'Two Dogmas of Empiricism'. The paper caused a kind of a revival of interest in Duhem's original formulation and it assumed a new life of the problem in a form of the so-called Duhem–Quine thesis.

The aim of the paper is to reconsider whether Duhem was right to argue that there are no crucial experiments in physics. In order to assess the validity of the thesis, I will (1) make an exposition of his arguments in its favour, and (2) analyse the major criticism of this position offered in the subject-literature by Adolf Grünbaum, who explicitly attacked arguments for the thesis as inconclusive and false. Finally, (3) I will try to present the possible ways of defence of the Duhem–Quine thesis and I will argue that the original formulation of the thesis is well qualified and plausible.

**David Teplow. Alzheimer's disease: Philosophical impediments towards a cure.** It is obvious that to study the history or philosophy of science, science itself must first exist. It is much less obvious to many, and especially to students,

that to 'do science' optimally requires an equally deep grounding in the history and philosophy of science. What applies to students applies no less, and possibly even more, to professors in the natural sciences. I discuss the clinical and scientific history of Alzheimer's disease, with special reference to controversies that have arisen from one of the most common and insidious errors of scientific practice, misassumption. Misassumptions will be exemplified through the consideration of a priori bias and inappropriate adherence to dogma. Examples of Kuhn-like paradigm shifts will be discussed. Concluding remarks address philosophical changes necessary if the Alzheimer's disease research community is to make progress towards a cure.

**Adam Toon. Where is the understanding?** There is now a growing interest in scientific understanding (e.g. de Regt et al. 2009). Understanding has often been felt to be too subjective to merit sustained discussion by philosophers of science. One reason for this is a tendency to identify understanding with the distinctive 'Aha!' feeling that we often experience when we explain something. However, recent work in both epistemology and philosophy of science has emphasised that, while it might be accompanied by a distinctive phenomenology, understanding is an important cognitive state that philosophers should seek to analyse.

Understanding poses a range of questions. For example, if understanding is a cognitive state, what is the nature of that state? Most authors agree that understanding goes beyond (mere) belief. In order to understand a phenomenon, a scientist must not only be able to recall relevant facts or theoretical principles; they must also 'grasp' or 'see' the connections between them. What are these acts of 'grasping' or 'seeing' (e.g. Grimm 2010)? And if understanding goes beyond simply believing, or even knowing, relevant facts and principles, how do explanations provide us with understanding (e.g. de Regt et al. 2009)?

In this paper, I will argue that scientific understanding is a form of situated cognition. Situated cognition is a growing movement in cognitive science that stresses the importance of interaction between brain, body and world in our cognitive processes (e.g. Robbins and Aydede 2009). Some authors have argued that such work provides a fruitful framework for studying scientific reasoning (e.g. Bechtel 1996, Giere 2006, Nersessian 2005). My aim in this paper will be to refine and develop a situated approach to scientific understanding and locate this approach within the recent discussion of understanding in epistemology and philosophy of science.

In doing so, I will argue the framework of situated cognition may be applied not only to explanatory inquiry or the act of giving an explanation, but also to understanding itself: the acts of 'grasping' or 'seeing' that many authors take to be characteristic of understanding often take place not in the scientist's head, but in processes that incorporate external, material devices. As well as defending this view against likely objections, I will consider its implications for a range of issues discussed within the recent literature on understanding, including the relationship between understanding and explanation.

Bechtel, W. (1996). What Should a Connectionist Philosophy of Science Look Like? In R. McCauley (Ed.) *The Churchlands and Their Critics*. Blackwell.

Giere, R. (2006) *Scientific Perspectivism*. Chicago University Press.

Grimm, S. (2010). The Goal of Explanation. *Studies in History and Philosophy of Science*. 41 (4): 337–344.

Nersessian, N. (2005). Interpreting scientific and engineering practices: Integrating the cognitive, social, and cultural dimensions. In M. Gorman, R. Tweney, D. Gooding, & A. Kincannon (Eds.), *Scientific and technological thinking* (pp. 17–56). Erlbaum.

Regt, H. de, Leonelli, S. and Eigner, K. (eds.) (2009). *Scientific Understanding: Philosophical Perspectives*. Pittsburgh University Press.

Robbins, P. and Aydede, M. (eds.) (2009). *The Cambridge Handbook of Situated Cognition*. Cambridge University Press.

### **Nick Tosh. Reviving finite frequentism: Humean chance without best systems.**

My analysandum is non-deterministic chance: roughly speaking, the kind that is often associated with quantum mechanics. My analytical strategy is finite frequentism: that is, I identify chances with relative frequencies of occurrence within actual, finite reference classes. Philosophers have long regarded this strategy as hopeless. I show that the standard objections become significantly less compelling if (i) we require reference classmates to have qualitatively identical histories; (ii) we assume relativistic (as opposed to Newtonian) causal structure; and (iii) we recognise that the actual world may, for all we know, be much larger than our own past light cone.

The version of frequentism I defend is metaphysically undemanding, makes no appeal to ‘objective’ measures of simplicity and informativeness, and recovers Lewis’s Principal Principle as a finitistic principle of self-location indifference. The advantages must be set against a couple of counterintuitive implications. According to my analysis, if chances (e.g. radium decay chances) are roughly as we take them to be, then the actual world contains enormously large collections of duplicate light cones. Furthermore, if the nomic facts about chance are roughly as we take them to be, then the statistics of these collections are tightly constrained by non-local laws. While surprising, these implications cannot be said to clash with current science. Indeed, recent work in the philosophy of cosmology has stressed the extent to which the global structure of our spacetime is underdetermined by observations we can (even in principle) make.

### **Dana Tulodziecki. The pessimistic meta-induction and the superfluity of approximate truth.**

The pessimistic meta-induction targets the realist’s claim that a theory’s (approximate) truth is the best explanation for its success. It attempts to do so by undercutting the alleged connection between truth and success by arguing that highly successful, yet wildly false theories are typical of the history of science, and, thus, that a theory’s success cannot be a symptom of its truth (cf. Laudan 1981, 1984).

There have been a number of prominent realist responses to the pessimistic meta-induction, most notably those of Worrall (1989), Kitcher, (1993), and Psillos (1999). All of these responses try to rehabilitate the connection between a theory’s (approximate) truth and its success by attempting to show that there is some kind of continuity between earlier and later theories, structural in the case of Worrall, and theoretical/referential in the cases of Kitcher and Psillos.

In this paper, I will argue that the extant realist responses to the pessimistic meta-induction are inadequate, since there are cases of theories that were both false and highly successful (even by the realist’s own, more stringent, criteria for success), but that, nevertheless, do not exhibit any of the continuities that have been suggested by realists as possible candidates for preservation.

I will make my case through discussing an example of such a theory: the 19th century miasma theory of disease. Specifically, I will show that this theory made a number of important and successful use-novel predictions, despite the fact that its central theoretical element – miasma – turned out not to exist. After showing that miasma was crucially involved in virtually every successful prediction the miasma theory made, I argue that not just is there no ontological continuity between the miasma theory and its successor, but neither can a case be made for any other kind of continuity, be it in terms of structure, laws, mechanisms, or kind-constitutive properties.

After discussing possible realist routes of escape from this predicament, and concluding that they all fail, I argue that the miasma case constitutes a new, particularly problematic, kind of counterexample to (all strands of) convergent realism, because the miasma theory’s successor – the germ theory – is not just taken to be approximately true, but completely true. Thus, not just does the miasma case show that there are cases of genuinely successful, yet false theories in which nothing gets retained (thus supporting the view that success is no sign of truth), but, more importantly, it also shows that truth can be achieved without any realist signs, specifically, that it can be gotten without the middle stage of approximate truth. The miasma case thus shows that approximate truth is not necessary for truth simpliciter, and that truth may be achieved without previous signs of convergence. By showing that approximate truth is superfluous, the case undermines the role that approximate truth is supposed to play for realists, thereby casting doubt on convergent realism itself.

### **Philippe Verreault-Julien. Understanding through counterfactual analysis modelling.**

I consider how some models, especially in economics, can yield understanding by way of counterfactual analysis. Models reputedly provide knowledge of the world because they represent it (Knuutila 2011). Moreover, the knowledge they yield is considered to be necessary to explain phenomena of interest. On most accounts, explanation requires faithful representation, that what explains be true (e.g. Hempel 1965; Salmon 1984; Woodward 2003; Craver 2006). But the conundrum is that some models hardly fulfil this requirement and nevertheless appear to be explanatory. In the context of economic modelling, Reiss (2012) has called this problem the ‘explanation paradox’. Economic models either faithfully represent, or they aren’t explanatory, or explanation doesn’t require faithful representation. But they can’t both misrepresent and be explanatory if explanations require a true explanans. Usual solutions consist in either, 1) showing that idealisations aren’t necessarily harmful and don’t prevent faithful representation of the explanatory information (e.g. Cartwright 1999; Mäki 2009; Hausman 2013), 2) denying altogether that models serve by themselves an explanatory function (e.g. Aydinonat 2007; Alexandrova 2008; Grüne-Yanoff 2009), or 3) amending traditional accounts of explanation so that faithful representation is no longer necessary (e.g. Cartwright 1983; Bokulich 2011; Kennedy 2012).

I argue for another option that might accord better with the prevalent accounts of causal explanations and with the fact that modelling appears to yield epistemic benefits in the form of understanding. I claim that some models provide counterfactual knowledge that contributes to our scientific understanding while not being explanatory. Explanatory understanding is itself often considered to be constituted by knowledge that allows to answer ‘what-if-things-had-been-different’ questions (w-questions) (e.g. Woodward 2003; Grimm 2006; Ylikoski

2009). Grasping certain counterfactual dependencies allows to correctly infer what would have happened were things different. Explanations thus contribute to understanding because they yield such knowledge. But explanations are arguably not the only route to counterfactual knowledge (Lipton 2009; Khalifa 2013).

Indeed, an underrated feature of non-representational models is that they do provide knowledge. Whereas they might not afford actual knowledge of target systems of interest, they may nevertheless afford counterfactual knowledge of causal or conceptual dependencies. This knowledge, while not explanatory in itself, might nevertheless contribute to our understanding since it allows to answer w-questions. Modelling can thus be seen as a kind of sophisticated analysis of counterfactual claims, which, for instance, are often deemed central in accounts of causation. These counterfactuals help to establish possible difference makers even though they might fall short of establishing actual causes. We know that in a given possible world some factor of interest depends on another, thus contributing to answer w-questions.

My account satisfies two desiderata others may not. First, we need not commit to problematic defences of idealisations or amend theories of explanation we regard as successful in order to account for the explanatoriness of misrepresenting models. Second, since counterfactual knowledge affords a kind of understanding that is similar to what explanations provide, we need not deny the epistemic benefits those models seem to provide despite the fact that they misrepresent.

**Charlotte Werndl. On defining climate and climate change.** The aim of the paper is to provide a clear and thorough conceptual analysis of the main candidates for a definition of climate and climate change. Of course, different definitions of climate and climate change are discussed in the climate science literature. However, what is missing is a clear and thorough conceptual analysis of the different definitions and their benefits and problems. This paper aims to contribute to filling this gap.

First climate variables and a simple example are introduced that will be used to illustrate the definitions of climate. Then five desiderata on a definition of climate are presented. A definition of climate should be empirically applicable (Desideratum 1), it should correctly classify different climates (Desideratum 2), it should not depend on our knowledge (Desideratum 3), it should be applicable to the past, present and future (Desideratum 4) and it should be mathematically well-defined (Desideratum 5).

With help of these desiderata the main definitions of climate are discussed. The first group of definitions discussed are distributions over time. According to the first definition, climate is the finite distribution of the climate variables over time for constant external conditions. This definition suffers the problem that it may be empirically void because in reality the external conditions are not constant (thereby violating Desideratum 1). With the second definition one tries to avoid this by defining climate as the finite distribution of the climate variables over time when the external conditions vary as in reality. However, this definition does not classify different climates correctly (thereby violating Desideratum 2).

According to the third definition, climate is the finite distribution of the climate variables over time relative to a regime of varying external conditions. This definition is novel and is introduced as a response to problems with Definition 1 and 2.

The second group of definitions discussed are ensemble distributions. According to the fourth definition, climate is the ensemble distribution of the climate variables for constant external conditions. However, this definition again suffers from the serious problem that it may be empirically void because the external conditions vary in reality (thereby violating Desideratum 1). With the fifth definition one tries to avoid this by defining climate as the ensemble distribution of the climate variables when the external conditions vary as in reality. While being attractive from a predictive perspective, this definition depends on our knowledge (thereby violating Desideratum 3), it is unclear how to define the past and present climate (thereby violating Desideratum 4) and there is no relation to the observational record (thereby violating Desideratum 1).

Infinite versions of Definitions 1-5 are also discussed. They are quickly dismissed since they suffer from the additional problems that they may be empirically void (thereby violating Desideratum 1) and that the relevant limits may not exist (thereby violating Desideratum 2).

The conclusion is that while the novel Definition 3 is promising because it satisfies all desiderata, the widely-endorsed Definitions 1, 2, 4 and 5 all suffer from serious problems.

**Adam White. Emergence in biological pathways.** I will argue that the claims made by Boogerd et al for emergence in biological pathways are unsubstantiated. Nevertheless I will suggest a plausible argument for why emergence in biological pathways might still occur.

Boogerd et al's paper analysis is based on C.D. Broad's theory of emergence. For Boogerd et al, a property is taken to be pathway emergent (henceforth: emergent) if it is a dynamic systemic property of a pathway that cannot be deduced, even in principle, from:

1. properties of that system's parts in isolation or in less complex systems and
2. proportions and organisation of the parts and
3. laws of composition.

Boogerd et al do not provide an argument for why pathways are sometimes emergent; instead they provide a case study which they claim demonstrates emergence.

The case study is *in silico* and consists of a simulation model of a hypothetical pathway of three reversible reaction steps:

$A \leftrightarrow B \leftrightarrow C \leftrightarrow D$  (Model M0)

A, B, C and D are reactants and there is a negative feedback loop. The dynamics of this pathway are compared with those of two simpler models:

$A \leftrightarrow B \leftrightarrow C$  (Model M1) and

$B \leftrightarrow C \leftrightarrow D$  (Model M2)

Models M0, M1 and M2 are all constructed using different combinations of just three rate laws and the 'kinetic law of composition' which combines rate laws. Boogerd et al show that M0 can have 'qualitatively different' dynamics from M1 and M2 and that the dynamics of M0 cannot be deduced from the systemic dynamics of M1 and M2.

However the case study does not demonstrate emergence. This is because the dynamics of M0 can be deduced from the dynamics of the individual reaction steps. Within biochemistry, the dynamics of a pathway are taken to be fully determined by the rate laws of the reaction steps within that pathway, initial conditions and the kinetic law of composition. The kinetic law of composition is taken as always being correct. Given this framework (and Boogerd et al's definition of emergence), a necessary

requirement for emergence is that a pathway has at least one rate law that is not present in simpler systems. This requirement is not satisfied in the case study.

Nevertheless there is a plausible argument for actual (as opposed to *in silico*) pathway dynamics sometimes being emergent. This is because rate constants in biochemistry are highly context sensitive. Many factors have been identified that contribute to this; e.g. there are often crowding and confinement effects. Rate laws are local fragile causal laws; they apply to a limited number of contexts and small changes in their context can 'break' a law. At present biologists are often not able to accurately deduce how rate constants will change with context. Perhaps this is sometimes indicative of pathway emergence.

Boogerd F.C., Bruggeman F.J., Richardson R.C., Stephan A. and Westerhoff H.V. 2005. Emergence and Its Place in Nature: A Case Study of Biochemical Networks. *Synthese* 145: 131–164.

### **Alastair Wilson. Towards a hybrid theory of laws.**

Recent work on laws of nature has centred around the dispute between 'Humeans' and 'anti-Humeans'. Do laws merely describe, or do they govern in some stronger sense? This paper explores the possibility of a middle way: a hybrid view which does justice both to the Humean's methodological motivations and to the anti-Humean's metaphysical motivations. The account treats law statements as context-sensitive; it implements this context-sensitivity via quantifier variance associated with quantification over qualitative features of modal space.

According to David Lewis, laws of nature are regularities entailed by whichever axiomatization of the facts about a world strikes the best balance between a) strength; and b) simplicity of expression in a fundamental language. Laws are compact statements of the occurrent facts: they describe rather than governing. This view has apparent epistemological and methodological advantages; it promises to make laws empirically respectable. But the view has costs; to many, it has seemed to strip the laws of an essentially modal element, and to limit their capacity to feature in explanations.

According to David Armstrong, laws of nature are relations amongst universals. They are contingent relations, since they could have been different; but they are relations of necessitation, since they guarantee that one property necessitates another. The modal force associated with these relations is said to support explanations which are unavailable to Humeans. But Lewis has presented an influential 'dormitive virtue' objection against this view, and it also faces a regress objection from Alexander Bird.

I argue that the dormitive virtue and regress objections to anti-Humeanism can be answered by adopting a strong form of necessitarianism about laws. This necessitarianism also permits of an unexpected unification of the Humean and anti-Humean viewpoints. Necessitarians can characterize the fundamental laws as those general principles which strike the best balance of simplicity and strength across the whole of modal space. We can then capture a wide range of more familiar systems of non-fundamental laws via restricted quantification over this space of possibilities; these laws will feature in the truth-conditions of law-statements in many ordinary contexts. We can also make a distinction between local and global laws, which has applications to the fine-tuning problem.

Depending on one's theory of modality, this approach promises all of the methodological virtues of the Humean approach. In the limiting case, modal realists can still retain the

central tenet of Humeanism unchanged: fundamental laws are simply the most general patterns in the totality of what exists.

My hybrid account most closely resembles one offered by Robert Pargetter in 1983. Pargetter argued that modal force could be added to Lewis' theory of laws by expanding the domain of the facts axiomatized by the laws, from facts about a single world to facts about a set of worlds inter-related by primitive external relations. I argue that Pargetter's theory is ultimately subject to the same objections as Armstrong's theory, because Pargetter also assumes the laws to be contingent.

### **Lena Zuchowski. Revisiting Smale's 14th problem: Are there two kinds of chaos?**

In 1998, Stephen Smale proposed the following question as one of eighteen mathematical problems to be solved in the 21st century:

'Is the dynamics of the ordinary differential equations of Lorenz that of the geometric Lorenz attractor of Williams, Guckenheimer, and Yorke?'

Later in the article, he clarifies that the question aims to establish whether the system of Lorenz equations is chaotic in the same sense as the horseshoe map he himself investigated in 1967. In 2002, Warwick Tucker showed that rigorously constructed numerical solutions to Lorenz's system support an attractor similar to the analytic one found by Williams, Guckenheimer and Yorke. For many, he had thereby solved Smale's 14th problem.

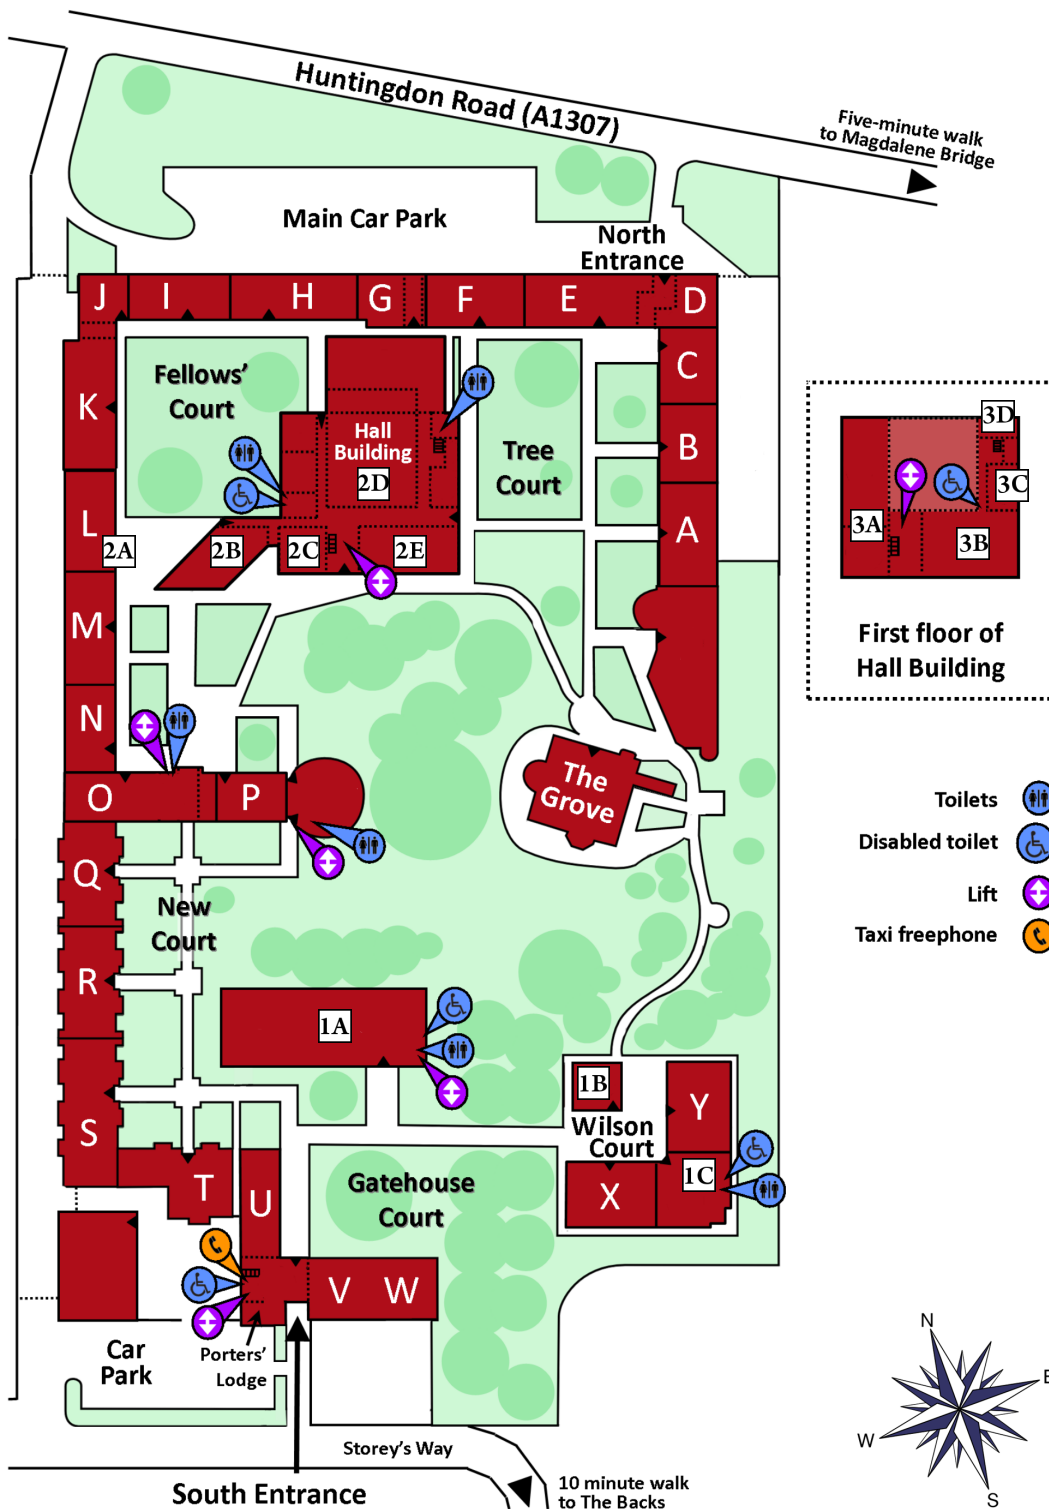
I will revisit Smale's 14th problem and address the question whether we can confidently class both properties of numerically integrated solutions as well as those of maps analytically constructed by subsequent self-application of a function as 'chaotic'. In the framework of philosophy of science, the question is less one of mere partial formal equivalence (as addressed by Tucker) but one of conceptual equivalence with respect to a given definition of 'chaos'.

I will argue that there are two prevalent classes of chaos definitions: one focusing on the existence of infinitely many periodic points and the other focusing on aperiodicity and pseudo-randomness. Furthermore, I will show that proponents of the first definition tend to build their definition from a catalogue of instances of self-applied maps while the latter definition is particularly applicable to numerically integrated functions. I will illustrate this difference on the example of the logistic equation, which can both be used to construct maps with infinitely many periodic points by successively applying it to itself as well as show a degree of aperiodicity when integrated numerically. I will also show that a third class of solutions to the logistic equation – analytically integrated ones – will not be chaotic under either of these definitions. This raises questions about the ontology of chaos.

Approaching the topic from the side of philosophy of science, my answer to the title question and Smale's 14th problem is hence a cautious negative: the Lorenz attractor is only 'chaotic' under a certain, restricted definition of 'chaos' and not in the same sense as Smale's horseshoe. As such, this case can be seen as an illustration of the general divide in chaos definitions (and catalogues of defining instances) I have argued for.







1A. Auditorium

1B. Gordon Cameron

1C. William Thatcher

2A. Trust

2B. Old SCR

2C. Walter Grave

2D. Hall

2E. Café/Bar

3A. Upper Halls 1 & 2

3B. Reddaway

3C. Gaskoin

3D. Music