

SPSP 2018

**The 7th biennial conference of the
Society for Philosophy of Science in
Practice**

**30 June – 2 July, 2018
Ghent University, Belgium**

<http://www.spsp2018.ugent.be/>

Contents

About SPSP	3
Organising Committees	5
Philosophy and Philosophy of Science at Ghent University	6
Practical Information	7
SPSP 2018 General Schedule	8
Abstracts of Plenary Lectures	10
Abstracts of Symposia	12
Abstracts of Contributed Papers	101

About SPSP

Philosophy of science has traditionally focused on the relation between scientific theories and the world, at the risk of disregarding scientific practice. In social studies of science and technology, the predominant tendency has been to pay attention to scientific practice and its relation to theories, sometimes wilfully disregarding the world except as a product of social construction. Both approaches have their merits, but they each offer only a limited view, neglecting some essential aspects of science. We advocate a philosophy of scientific practice, based on an analytic framework that takes into consideration theory, practice and the world simultaneously.

The direction of philosophy of science we advocate is not entirely new: naturalistic philosophy of science, in concert with philosophical history of science, has often emphasized the need to study scientific practices; doctrines such as Hacking's "experimental realism" have viewed active intervention as the surest path to the knowledge of the world; pragmatists, operationalists and late-Wittgensteinians have attempted to ground truth and meaning in practices. Nonetheless, the concern with practice has always been somewhat outside the mainstream of English-language philosophy of science. We aim to change this situation, through a conscious and organized programme of detailed and systematic study of scientific practice that does not dispense with concerns about truth and rationality.

Practice consists of organized or regulated activities aimed at the achievement of certain goals. Therefore, the epistemology of practice must elucidate what kinds of activities are required in generating knowledge. Traditional debates in epistemology (concerning truth, fact, belief, certainty, observation, explanation, justification, evidence, etc.) may be re-framed with benefit in terms of activities. In a similar vein, practice-based treatments will also shed further light on questions about models, measurement, experimentation, etc., which have arisen with prominence in recent decades from considerations of actual scientific work.

There are some salient aspects of our general approach that are worth highlighting here.

(1) We are concerned with not only the acquisition and validation of knowledge, but its use. Our concern is not only about how pre-existing knowledge gets applied to practical ends, but also about how knowledge itself is fundamentally shaped by its intended use. We aim to build meaningful bridges between the philosophy of science and the newer fields of philosophy of technology and philosophy of medicine; we also hope to provide fresh perspectives for the latter fields.

(2) We emphasize how human artifacts, such as conceptual models and laboratory instruments, mediate between theories and the world. We seek to elucidate the role that these artifacts play in the shaping of scientific practice.

(3) Our view of scientific practice must not be distorted by lopsided attention to certain areas of science. The traditional focus on fundamental physics, as well as the more recent focus on certain areas of biology, will be supplemented by attention to other fields such as economics and other social/human sciences, the engineering sciences, and the medical sciences, as well as relatively neglected areas within biology, physics, and other physical sciences.

(4) In our methodology, it is crucial to have a productive interaction between philosophical reasoning and a study of actual scientific practices, past and present. This provides a strong rationale for history-and-philosophy of science as an integrated discipline, and also for inviting the participation of practicing scientists, engineers and policymakers.

SPSP 2018 is financially supported by:

- Research Foundation – Flanders (FWO).
- the research fund of the Faculty of Arts and Philosophy of Ghent University.

Organising Committees

SPSP steering committee

Chiara Ambrosio, University College London

Rachel Ankeny, University of Adelaide

Justin Biddle, Georgia Institute of Technology

Till Grüne-Yanoff, Royal Institute of Technology (KTH), Stockholm

Sabina Leonelli, University of Exeter

Matthew Lund, Rowan University

Joseph Rouse, Wesleyan University

Organising committee 7th biennial conference

Chiara Ambrosio, University College London

Rachel Ankeny, University of Adelaide

Justin Biddle, Georgia Institute of Technology

Leen De Vreese, Ghent University

Steffen Ducheyne, Free University of Brussels (VUB).

Raoul Gervais, University of Antwerp

Till Grüne-Yanoff, Royal Institute of Technology (KTH), Stockholm

Jan Heylen, KU Leuven

Sabina Leonelli, University of Exeter

Bert Leuridan, University of Antwerp

Matthew Lund, Rowan University

Joseph Rouse, Wesleyan University

Jean Paul van Bendegem, Free University of Brussels (VUB)

Maarten Van Dyck, Ghent University

Dingmar van Eck, Ghent University

Bart Van Kerkhove, Free University of Brussels (VUB) & Hasselt University

Erik Weber, Ghent University

Sylvia Wenmackers, KU Leuven

Local organising team UGent

Erik Weber (chair), Mathieu Beirlaen, Inge De Bal, Roxan Degeyter, Rogier De Langhe, Leen De Vreese, Federico Faroldi, Stef Frijters, Julie Mennes, Joke Meheus, Pawel Pawlowski, Dingmar van Eck, Frederik Van de Putte & Dietlinde Wouters.

Philosophy and Philosophy of Science at Ghent University

The Centre for Logic and Philosophy of Science which organises SPPS2018 is part of the Department of Philosophy and Moral Science. Our department offers a bachelor and master programme in philosophy, and a bachelor and master programme in moral science. Our philosophy programme covers the traditional topics: history of philosophy from ancient to contemporary philosophy, epistemology, logic, philosophy of science, metaphysics, philosophical anthropology and theoretical and applied ethics. The aim is to give our students an advanced knowledge and grasp of theories, methods and skills in these fields. Our programme in moral science has a different focus: it contains less logic, epistemology, philosophy of science and history of philosophy. Students in moral science are trained in empirical research methods, which allow them to study moral phenomena in a descriptive way (as opposed to the normative approach in philosophical ethics) and get a substantial background in the social sciences and psychology.

The Centre for Logic and Philosophy of Science was founded in 1993. Most of the research that is done at the centre fits into three research lines:

- Logical analysis of scientific reasoning processes
- Methodological and epistemological analysis of scientific reasoning processes, and
- Integrated history and philosophy of science

Examples of specific topics that fit into the first research line are: logical analyses of paraconsistent reasoning, reasoning under uncertainty, defeasible reasoning, abduction, causal reasoning, induction, analogical reasoning, belief revision, theory change, conceptual change, etc.

Examples of specific topics that fit into the second research line are: methodological and epistemological analyses of causation and mechanisms, scientific discovery, the structure of scientific theories and models, experiments and thought experiments, theory choice, theory dynamics, rationality, etc.

The research in integrated history and philosophy of science includes work on scientists and philosophers such as Descartes, Euler, Galilei, Leibniz, Mach, Maxwell, Newton, Poincaré, and Stevin.

Practical Information

Coffee and lunch

Coffee, tea and lunch are included in the registration fee and will be served during the breaks. Vegetarian and vegan lunches will be served together with the regular lunch buffet so please pay attention to the sign at the buffet.

Book exhibit

The book exhibit is in the registration room (120.015).

Internet

Make a wireless connection with "UGentGuest". If you have set up to request an IP address automatically, you will receive an IP address starting with 193.190.8x.

Now you are connected, but not yet authenticated. You should start a webbrowser and you will be redirected to a logon screen. If not surf to <http://www.ugent.be>. Enter the username and password:

- Login: guestSsp20
- Password: pDYZtHNU

After correct authentication you can use the Internet connection. Your connection to this wireless LAN is not encrypted. To protect your personal data, please use encrypted connections like https, imaps, ssh etc. or a VPN client.

You are not allowed to pass on the login information to others.

Detailed programme

Detailed up-to-date programmes are available at the registration desk and on the website (<http://www.spsp2018.ugent.be/programme/>)

Twitter

hashtag: #SPSP2018

SPSP 2018 General Schedule

FRIDAY 29 JUNE

9:45 – 17:00	Pre-conference workshop
18:00 – 20:00	Registration Pre-conference informal gathering

SATURDAY 30 JUNE

8:15 – 9:00	Registration
8:50 – 9:00	Opening session: Welcome
9:00 – 10:10	Plenary Lecture 1
10:10 – 10:30	Coffee Break
10:30 – 12:00	Concurrent Sessions 1A – 1G
12:00 – 13:30	Lunch Poster session
13:30 – 15:30	Concurrent Sessions 2A – 2F
15:30 – 16:00	Coffee Break
16:00 – 17:30	Concurrent Sessions 3A – 3F

SUNDAY 1 JULY

9:00 – 11:00	Concurrent Sessions 4A – 4F
11:00 – 11:30	Coffee Break Poster session
11:30 – 12:40	Plenary Lecture 2
12:40 – 14:00	Lunch Philosophy of Medicine Networking Lunch
14:00 – 16:00	Concurrent Sessions 5A – 5F
16:00 – 16:30	Coffee Break Poster session
16:30 – 17:40	Plenary Lecture 3
19:30 – 22:00	Conference dinner

MONDAY 2 JULY

9:00 – 11:00	Concurrent Sessions 6A – 6F
11:00 – 11:30	Coffee Break
11:30 – 12:40	Plenary Lecture 4
12:40 – 14:00	Lunch Newsletter meeting
14:00 – 15:30	Concurrent Sessions 7A – 7F
15:45 – 16:30	Closing session

SPSP 2018 Abstracts

Plenary Lectures (in order of appearance)

Visualization Practices

William Bechtel – University of California, United States

From (1) the formulation of a hypothesis to (2) the design of instruments and experiments, (3) the representation of data, and (4) the formulation of the knowledge acquired, scientists rely on visualizations. Whereas philosophers often confine themselves to one-dimensional linguistic representations, scientists use linguistic representations as supplements or guides to visual representations. In part they prefer visual representations because they enhance the dimensionality of the representations (first to two, but through shape, color, etc., to many more). I will begin with a brief overview of visual representations used to formulate hypotheses, represent experimental instruments and protocols, present data, and communicate conclusions. I will then focus in on two specific visualization practices—the growing practice of publishing visual abstracts and the development of platforms to represent and analyze networks.

Philosophy of Science in the Age of Big and Open Data: What Does “Studying Practice” Involve?

Sabina Leonelli – University of Exeter, United Kingdom

I examine three modes of analysis that I see as characteristic – particularly when they are intertwined – of philosophy of science in practice: (1) the study of science in the world, which identifies concerns and features of research endeavors (past or present) in ways that are open to empirical scrutiny; (2) the development of philosophy relating to such study, which provides ways to critically reflect upon and situate the practices and outputs of research vis-à-vis their epistemological, social, ethical and/or ontological significance; and (3) the engagement of society in such philosophical ideas, where “society” can include any type and number of institutions and groups which are affected by research processes (including of course philosophers themselves). I illustrate these modes of analysis through my own work on the philosophy of big and open data.

Did Galileo's Experiments Confirm the Law of Fall?

Maarten Van Dyck – Ghent University, Belgium

Alexandre Koyré's *Etudes Galiléennes*, first published in 1939, started a long-lasting debate on the question whether Galileo actually executed experiments with bodies descending from an inclined plane, and in case he did, whether their results could confirm his law of fall. The long life of the debate can be explained by its perceived role in setting the agenda for interpretations of the scientific revolution. After its initial stage, which was focused on philosophical questions, the debate was shaped by the use of material reconstructions and the attempt to decipher new manuscript evidence. Historically there is no doubt anymore that Galileo engaged in a sustained program of experimental activity, but the impact of these experiments remains an open question.

In my talk I will start by reassessing Koyré's claims, which are much more subtle than is usually thought. I will then sketch the development of Galileo's thinking, and the crucial point at which experiments started playing a new role. This will show the complex web of entangled factors which needed to be put into play before it made sense to start interpreting the results of inclined plane experiments as confirming the law of fall. I will end by using this to reflect on both the abstract and a-historical notion of "confirmation" and the general and totalizing category of the "scientific revolution".

Philosophy in the Field: Witnessing and Translating

Alison Wylie – University of British Columbia, Canada

The Canadian Truth and Reconciliation Commission (2015) calls on non-Indigenous Canadians to build equitable, respectful and transparent partnerships with Indigenous Peoples as the primary means for advancing reconciliation. In this spirit a UBC-based research cluster, Indigenous/Science, is building collaborative partnerships designed to bring the tools of archaeological science to bear on Indigenous-led research questions in a way that embodies a "practice of reconciliation." Integral to the work of this cluster is a commitment to document and assess the process of building these partnerships; a "reflection" working group is tasked with asking how collaborative partners can best navigate asymmetries of power and hierarchies of expertise as well as the cultural differences in ethical/epistemic stance that structure collaborative projects. I trace the background to this initiative in terms of several key junctures in my own turn to practice, and explore the question of what we philosophers can contribute to such ventures.

Symposia (alphabetical by last name of organizer)

Symposium: Virtue Epistemology of Mathematical Practice

Organizer: Andrew Aberdein.

Contributors:

Andrew Aberdein.

Colin Rittberg.

Fenner Tanswell.

Twenty-first century philosophy of mathematics has exhibited a turn towards issues of mathematical practice [4]. This represents a shift from the field's traditional focus on metaphysical questions to a focus on epistemology, in particular social epistemology. One of the most productive programmes in twenty-first century epistemology has been virtue epistemology. These parallel developments suggest that the time is ripe for a virtue epistemology of mathematics.

George Pólya, one of the most influential precursors to the practice turn in philosophy of mathematics, identified 'intellectual courage, intellectual honesty, and wise restraint' as the 'moral qualities of the scientist' that he considered indispensable [6]. However, these virtues are seldom explicitly invoked in his work or that of his followers. Indeed, until recently, discussion of virtues in the philosophy of mathematics has been fleeting and fragmentary at best. In the last few years this has begun to change. Epistemic virtues have attracted attention in the philosophy of science as components of a full account of successful theory choice [5], an argument that readily extends to mathematics. Within the philosophy of mathematical practice itself, attention to virtues has emerged from a variety of disparate sources. For example, Penelope Maddy's work on the justification of axioms appeals to theoretical virtues [3], an aspect of her work developed further by others [7]; Noel Clemente has suggested that aprioricity of axioms could be analysed in terms of a reliabilist epistemic virtue [2]; Fenner Tanswell has proposed virtue epistemology as the correct epistemology for mathematics (and perhaps even as the basis for progress in the metaphysics of mathematics) [9]; Francis Su has advocated an understanding of social utility of mathematical practice grounded in virtue ethics [8]; and Don Berry has put forward an account of the

role of proof in terms of theoretical virtues of permanence, reliability, autonomy, and consensus [1].

This panel brings together several researchers who have begun to study mathematical practice from the perspective of virtue epistemology with the intention of consolidating and encouraging this trend.

[1] Berry, D. (2018). Proof and the virtues of shared enquiry. *Philosophia Mathematica*. Forthcoming.

[2] Clemente, N. L. (2016). A virtue-based defense of mathematical apriorism. *Axiomathes*, 26:71–87.

[3] Maddy, P. (2011). *Defending the Axioms: On the Philosophical Foundations of Set Theory*. Oxford University Press, Oxford.

[4] Mancosu, P., ed. (2008). *The Philosophy of Mathematical Practice*. Oxford University Press, Oxford.

[5] Paternotte, C. and Ivanova, M. (2017). Virtues and vices in scientific practice. *Synthese*, 194:1787–1807.

[6] Pólya, G. (1954). *Mathematics and Plausible Reasoning*. Two Volumes. Princeton University Press, Princeton, N.J.

[7] Rittberg, C. and Tanswell, F. (2016). Mathematical values and set-theoretic pluralism. Presented at Set Theoretic Pluralism Symposium, Aberdeen.

[8] Su, F. E. (2017). Mathematics for human flourishing. *The American Mathematical Monthly*, 124(6):483–493.

[9] Tanswell, F. (2016). *Proof, Rigour & Informality: A Virtue Account of Mathematical Knowledge*. PhD thesis, University of St. Andrews.

Virtues, arguments, and mathematical practice

Andrew Aberdein – Florida Institute of Technology, United States

Several authors have proposed argumentation theory as a methodology for the study of mathematical practice [2]. Formal logic serves the traditional purposes of philosophy of mathematics very well. However, the philosophy of mathematical practice is concerned not just with formal derivation but with the social processes whereby mathematicians gain assent for their conjectures. Since formal logic is less well-adapted to the analysis of arguments in this wider sense, it is natural to look beyond it to argumentation theory, a discipline concerned with the analysis of natural language argument. Several authors have proposed virtue theory as an approach to argumentation theory [1]. Virtue theories of argument shift the focus away from arguments as abstractions onto the interpersonal nature of argumentation, stressing the

importance of arguers, respondents, and audiences, and especially the character of these participants.

Despite some overlap amongst their advocates, these two trends have never been addressed together. In doing so, it is natural to ask if their conjunction entails a virtue theoretic approach to mathematical practice: must the virtue theorist of argument also be a virtue theorist of mathematical practice? A negative answer to this question is not impossible. It could be held that those aspects of mathematical practice that lend themselves best to analysis in terms of argument do not correspond to features of argumentation theory where a virtue approach is of most value. In particular, some virtue theorists of argument deny that theirs is an all-embracing account, insisting that some issues, notably the appraisal of arguments, must be handed over to another theory [3].

Nonetheless, this paper defends a virtue argumentation theory of mathematical practice. It does so on two grounds. Firstly, there are significant but neglected areas of both argumentation theory and the study of mathematical practice where a shared virtue approach is potentially salutary. For example, conventional approaches in each discipline pay little attention to the contribution the respective practice makes to human flourishing [4]. Secondly, mathematical practice is potentially a valuable testbed for the ambitious varieties of virtue argumentation theory. Virtue accounts have already been proposed for aspects of mathematical practice corresponding to argument appraisal, such as the social acceptance of proofs. The success of such accounts would suggest that virtue approaches can be of comparable utility within argumentation in general.

[1] Aberdein, A. and Cohen, D. H. (2016). Introduction: Virtues and arguments. *Topoi*, 35(2):339–343.

[2] Aberdein, A. and Dove, I. J., eds (2013). *The Argument of Mathematics*. Springer, Dordrecht.

[3] Gascón, J. Á. (2016). Virtue and arguers. *Topoi*, 35(2):441–450.

[4] Su, F. E. (2017). Mathematics for human flourishing. *The American Mathematical Monthly*, 124(6):483–493.

[5] Tanswell, F. (2016). *Proof, Rigour & Informality: A Virtue Account of Mathematical Knowledge*. PhD thesis, University of St. Andrews.

Intellectual Humility in Mathematics

Colin Rittberg – Vrije Universiteit Brussel, Belgium

Intellectual Humility has to do with a virtuous understanding of the confidence we have for a belief (Kidd 2016) and the “owning” of our epistemic limitations (Whitcomb et al 2017). The demand for proofs in mathematics may thus look like a mechanism which guides mathematicians towards being intellectually humble. In this talk, I will explore the notion of Intellectual Humility in the context of mathematical practices with an eye to further our understanding of what good mathematics is beyond the right and wrong of mere calculations.

Kidd, Ian James. "Intellectual humility, confidence, and argumentation." *Topoi* 35, no. 2 (2016): 395-402.

Whitcomb, Dennis, Heather Battaly, Jason Baehr, and Daniel Howard-Snyder. "Intellectual humility: Owning our limitations." *Philosophy and Phenomenological Research* 94, no. 3 (2017): 509-539.

Proof, Rigour and Mathematical Virtues

Fenner Tanswell – University of St Andrews, United Kingdom

In this paper I propose the application of virtue epistemology to mathematical knowledge. I argue that this provides us with the tools to account for informal proofs and the nature of rigour as they are found in mathematical practices, overcoming obstacles that rule out the opposing formalist-reductionist approach. Furthermore, virtue-theoretic terminology allows us to make sense of a great deal of other phenomena of mathematics in practice.

Symposium: Mathematical Proofs in Mathematical Practice

Organizers: Line Andersen, Joachim Frans, Yacin Hamami.

Contributors:

Line Andersen, Henrik Kragh Sørensen & Mikkel Willum Johansen.

Yamin Hamami & Rebecca Morris.

Joachim Frans.

Forty years ago, mathematician and philosopher Reuben Hersh suggested that we reject formalism as a philosophy of mathematics; that we give up the picture of mathematics as formal derivations from some given set of formulas.

Hersh wanted to replace formalism with “a philosophy that is true to the reality of mathematical experience.” (1979, p. 40) Since then many philosophers have used case studies to examine how ordinary mathematical proofs, as they occur in mathematical practice, differ from formal derivations (e.g., Lakatos 1974; Rav 1999; Manders 2008). Larvor (2012) goes further and proposes a general account of informal proof on which informal inferences are conceived as actions on propositions and objects (such as diagrams and notational expressions). In this session, we continue this general line of work. Like Larvor, we explore the benefits of focusing on the proving agents when thinking about mathematical proofs. Furthermore, we examine the fruitfulness of studying proofs in the context of their audiences. Thus, this session attempts to shed new light on the nature of ordinary mathematical proofs by looking at them in the context of the proving agents and the audiences. In the first talk, it is argued that the rationality of deductive inferences in proofs ought to be conceived in terms of the rationality of deductive inferences in proving activities. A deductive inference in a proving activity is, in turn, rational whenever it forms part of a plan that accords with the norms of rationality of planning agency. A conception of the planning agency underlying proving activities is developed. The second talk also focuses on the proving agents, but asks how they take into account the intended audience of the proof when developing it. The talk is based on interviews with two mathematicians about how they prepared the proofs in a joint article for validation by the reader. Like the second talk, the third talk considers proofs in the context of their audience, but this talk focuses on the explanatoriness of proofs. It is argued that whether a proof explains a theorem, as opposed to just establishes that the theorem is true, depends on the context; it depends on the audience and on the kind of activity in which the proof is used.

Hersh, R. (1979). Some proposals for reviving the philosophy of mathematics. *Advances in Mathematics*, 31, 31–50.

Lakatos, I. (1976). *Proofs and refutations: The logic of mathematical discovery*. New York: Cambridge University Press.

Larvor, B. (2012). How to think about informal proofs. *Synthese*, 187, 715–730.

Manders, K. (2008). The Euclidean diagram. In P. Mancosu (Ed.) (2008). *The philosophy of mathematical practice* (pp. 80–133). Oxford et al.: Oxford University Press.

Rav, Y. (1999). Why do we prove theorems? *Philosophia Mathematica*, 7, 5–41.

Mathematicians Writing for Mathematicians: The Framing of Proofs

Line Andersen – Aarhus University, Denmark

Henrik Kragh Sørensen – University of Copenhagen, Denmark

Mikkel Willum Johansen – University of Copenhagen, Denmark

As mathematician William Thurston has pointed out, mathematicians “prove things in a certain context and address them to a certain audience.” (1994, 175) Ordinary mathematical proofs, as found in mathematical practice, appeal to the intuitions and background knowledge of the intended reader. Hence, what such proofs look like depends on their intended audience, and to understand the nature of mathematical proofs as presented, we need to examine how they are made to address their audience. Some studies take into account the intended audience of mathematical proofs when discussing the level of granularity of proofs; they point out that the appropriate level of granularity of a proof depends on the audience (e.g., Fallis 2003; Larvor 2012; and Paseau 2016). For example, the level of granularity of a textbook proof written for high school students will often be higher than the level of granularity of research proofs. However, these studies do not go into detail with how the level of granularity of a proof is made to fit the audience. In this talk, we focus on how the level of granularity of a research proof is made to fit the intended expert audience and, more generally, on how mathematicians frame their proofs when writing for mathematicians. We have conducted interviews with two mathematicians, the talented PhD student Adam and his experienced supervisor Thomas, about a research article they wrote together. Over the course of two years, Thomas and Adam revised Adam’s very detailed first draft. At the beginning of this collaboration, Adam was a new PhD student and did not know how to write for mathematicians, but he was very knowledgeable about the subject of the article. Thus, one main purpose of revising the article was to make it take into account the intended audience. For this reason, the changes made to the initial draft and the authors’ purpose in making them provide a window to how mathematicians write for mathematicians. We examined how their article prepares their proofs for validation by the reader and found that it prepares the proofs for two types of validation that the reader can easily switch between: line-by-line validation and another type of validation which we will describe in detail in the talk. The two types of validation do not require the same level of granularity of the proofs, and the proofs are thus made to present themselves to the reader at two different levels of granularity.

Fallis, D. (2003). Intentional gaps in mathematical proofs. *Synthese*, 134, 45–

69.

Larvor, B. (2012). How to think about informal proofs. *Synthese*, 187, 715–730.
Paseau, A. C. (2016). What’s the point of complete rigour? *Mind*, 125, 177–207.
Thurston, W. P. (1994). On proof and progress in mathematics. *Bulletin of the American Mathematical Society*, 30, 161–177.

Rationality in Mathematical Proofs

Yacin Hamami – Vrije Universiteit Brussel, Belgium

Rebecca Morris – Stanford University, United States

On the traditional view, a mathematical proof is nothing more than a sequence of deductive steps, the only requirement being that each deductive step be valid. Although the traditional view is perfectly adequate if one is primarily concerned with the capacity of mathematical proofs to provide justification for mathematical knowledge, it has been argued by several leading mathematicians that it offers a too impoverished conception of the nature of mathematical proofs. Henri Poincaré is a case in point when he writes: “A mathematical demonstration is not a simple juxtaposition of syllogisms; it consists of syllogisms placed in a certain order, and the order in which these elements are placed is much more important than the elements themselves.” (Poincaré, 1908, p.49)

What Poincaré is contesting here is the idea that mathematical proofs are sequences of arbitrary deductive steps. He is thus promoting a view according to which a mathematical proof is a sequence of *rational* deductive steps. Fleshing out this view requires then an account of what it means for a deductive step in a mathematical proof to qualify as rational. Providing such an account calls for a radical shift from the static, agent-free perspective on mathematical proofs inherent to the traditional view. Our starting point is the observation that the notions of deductive step and mathematical proof have direct counterparts in the realm of action. Corresponding to the notion of deductive step is the notion of deductive inference, a deductive inference being first and foremost an action of an epistemic nature (Prawitz, 2012; Boghossian, 2014). As a counterpart to the notion of mathematical proof, we shall introduce the notion of a proving activity, a proving activity being defined as a sequence of rational deductive inferences. Now, our claim is that the rationality of deductive steps in mathematical proofs is derivative from the rationality of deductive inferences in proving activities, and that an account of the former ought to be derived from an account of the latter.

In order to articulate such an account, we will first observe that proving activities share two essential features with many of our ordinary human activities: they are goal-directed and temporally extended activities. As was repeatedly emphasized by Michael Bratman (1987, 2007), for cognitively limited agents like us, the realization of such activities requires a form of planning agency. The core of our account will then consist in developing a conception of the planning agency underlying proving activities, with a particular attention to the practical reasoning central to it. We will then reach an account of the rationality of deductive inferences in proving activities by saying that a deductive step in a proving activity is rational whenever it figures in a plan that has been elaborated according to the norms of rationality of planning agency. Our whole discussion will be illustrated by a concrete example of an ordinary mathematical proof in mathematical practice. We will conclude the talk by pointing out the potential interests of conceiving mathematical proofs as sequences of rational deductive steps—that is as products of rational planning agents—for several issues currently on the research agenda of the philosophy of mathematical practice (Mancosu, 2008).

Paul Boghossian. “What is inference?” *Philosophical Studies*, 169(1):1–18, 2014.

Michael E. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.

Michael E. Bratman. *Structures of Agency: Essays*. Oxford University Press, Oxford, 2007.

Paolo Mancosu. *The Philosophy of Mathematical Practice*. Oxford University Press, Oxford, 2008.

Henri Poincaré. *Science et Méthode*. Flammarion, Paris, 1908.

The objectivity/subjectivity of the distinction between explanatory and non-explanatory proofs

Joachim Frans – Vrije Universiteit Brussel, Belgium

Many mathematicians and philosophers of mathematics agree that proofs are central to mathematics. The reason for this pivotal role should not, however, be reduced to their role in establishing the truth of a mathematical claims. Similar to other scientists, activities of mathematicians are not merely guided by the aim of accumulating truths. Another aim, for example, might be to provide explanations of those mathematical truths. In recent decades, the particular notion of an explanatory proof has caught the interest of philosophers. While all proofs show that a theorem is true, an explanatory

proof goes further and also reveals why the theorem is true. The most well-known proposal to explicate the nature of an explanatory proof is presented by Steiner (1978). Steiner introduces the notion of a characterizing property between explanatory and non-explanatory proofs. Although considered to be an interesting starting point, this model is not free of criticisms. A specific critique I am interested in for this talk comes from Resnik and Kushner (1987). They argue that Steiner's objective to objectively capture a distinction between explanatory and non-explanatory proofs is misguided, since they believe the explanatoriness of a proof depends on specific contexts. This argument connects closely to Resnik and Kushner's endorsement of van Fraassen's (1980) pragmatic account of scientific explanation. Sandborg (1998), on the other hand, has argued that the application of van Fraassen's account to mathematics is problematic. Turning to more recent contributions in the literature concerning explanatory proofs, we see that several philosophers have proposed alternatives to Steiner's model. Nonetheless, critical remarks on conceiving the distinction between explanatory and non-explanatory as objective can still be found. Zelcer (2013) argues that there is no explanation in mathematics, at least not in an objective sense. So the models of explanatory proof provided by Steiner and other philosophers are meaningless. Inglis and Aberdein (2016) argue that a consensus between mathematicians which proofs should be labelled as explanatory is lacking. Consequently, models that are developed by means of the intuition that a certain proof is clearly explanatory are problematic from the start.

I will argue that there is much to say for the claim that the notion of an explanatory proof cannot be fleshed out in purely objective terms. This should not necessarily be seen as a fundamental problem for philosophical investigations into this notion. Rather, I propose a change of perspective. Instead of emphasizing the role of the proof itself, I suggest to look at activities related to doing something with the proof, and how they connect to the aim of giving mathematical explanations. From this perspective, current models of explanatory proofs can be re-evaluated and be used to make context-dependent claims. Moreover, this can lead to embracing contextual aspects of mathematical explanation that go beyond the problematic application to mathematics of van Fraassen's pragmatic account.

Inglis, M. & Aberdein, A. (2016) Diversity in proof appraisal. In: Larvor, B. (ed.) *Mathematical Cultures*, Birkhäuser: pp. 163-179.

Resnik, M.D. & Kushner, D. (1987). Explanation, independence and realism in mathematics. *The British journal for the philosophy of science*, 38(2):pp. 141-158.

Sandborg, D. (1998) *Mathematical explanation and the theory of why-*

questions. *The British Journal for philosophy of science*. 49(4): pp. 603-624.
Steiner, M. (1978) Mathematical explanation. *Philosophical studies*, 34(2): pp. 135 - 151.

Van Fraassen, B.C. (1980) *The scientific image*. Oxford, Oxford University Press.

Zelcer, M. (2013) Against mathematical explanation. *Journal for general philosophy of science*. 44(1): pp. 173-192.

Symposium: Philosophical Issues in Cancer Diagnosis and Treatment

Organizer: Justin Biddle.

Contributors:

Justin Biddle.

Bennett Holman.

Stephen John.

Recent discussions of “overdiagnosis” have highlighted the importance of ethical and epistemological considerations in medical diagnosis and treatment, particularly in early detection programs of cancer. The issues raised here fall at the intersection of a number of different fields, including bioethics, philosophy of medicine, clinical practice, and philosophy of science in practice. This symposium will explore a number of these issues, particularly as they relate to philosophy of science in practice.

Traditional approaches to early diagnosis and treatment of cancer emphasize prevention; within these approaches, it is assumed that early detection is best, as this allows patients to undergo treatment and prevent symptoms or death that would otherwise occur. Recently, however, concerns have been raised about early detection. It is now understood that some cancers either do not progress or progress so slowly that they will never cause symptoms or death, and many early detection programs cannot distinguish these cancers from those that will advance to cause symptoms. Early detection can thus lead to patients undergoing treatment unnecessarily and suffering consequent complications. There are challenging philosophical issues raised by early detection programs, including epistemological issues in diagnosis, ethical questions about physician-patient relationships, and how medical diagnosis should be modeled in the first place. The three presenters in this symposium will examine these issues.

Justin Biddle will examine the processes of risk assessment in prostate cancer screening. Drawing upon recent work on inductive and epistemic risk, he argues that prostate cancer diagnosis often involves value-laden judgment calls

on the part of physicians and that the values underlying these judgments often go unnoticed by the physicians making them. This fact poses challenges for effective risk communication to patients, and it has important implications for the ethical norms that should govern guidelines for clinical practice. Bennett Holman will examine the use of artificial intelligence (AI) in cancer diagnosis and prognosis. He reviews the accomplishments of AI in diagnosis and discusses the prospects for using AI to automate treatment recommendations for patients. He argues that properly implemented deep learning could solve a number of problems with both the technical and the social hindrances to reliable medical judgments. Finally, Stephen John draws on Hacking's work on representing and intervening to examine how we should conceptualize the act of medical diagnosis. He sets out a model of diagnosis as a particular kind of speech act, an exercitive, which involves changing the status of some individual within the healthcare system, and he develops the implications of this argument for accounts of disease and for the role of "non-epistemic values" in scientific justification.

Epistemic Risks in Prostate Cancer Diagnosis: Implications for Ethics and Policy

Justin Biddle – Georgia Institute of Technology, United States

Cancer screening—or testing for cancer in the absence of symptoms—is the subject of much debate. Screening has the potential to save lives by identifying and treating cancers in early stages. However, not all cancers cause symptoms, and the diagnosis of these cancers can lead to unnecessary treatments and subsequent side-effects and complications. The debate over cancer screening is a part of a larger discussion about overdiagnosis of disease, and at the very least, the debate has highlighted the difficulties involved in balancing the risks of failing to treat against the risks of overdiagnosis (e.g., Welch et al. 2011). This paper will focus on the debate over prostate cancer screening. In 2016, the United States Preventative Services Task Force (USPSTF) published draft recommendations that men ages 55 to 69 “make an individualized decision about prostate cancer screening with their clinician” (USPSTF 2016). This recommendation is updated from 2012 guidelines recommending that no man of any age undergo screening. This shift in recommendations, which opens the door to more men getting screened, is based on a shift in values. The 2012 guidelines are paternalistic and based on the norm of beneficence; the USPSTF determined on the basis of a risk-benefit calculation that it was not in the best interests of men to undergo screening, and as such, it recommended against screening. The 2016 recommendations, by contrast, are

based on the norm of respect for patient autonomy; they attempt to ensure that “each man has an opportunity to understand the potential benefits and harms of screening and to incorporate his values and preferences into his decision” (USPSTF 2016).

Respect for patient autonomy requires, at a minimum, that doctors communicate clearly to patients the risks and benefits of treatment options. Drawing upon recent work on inductive and epistemic risk, I examine the processes of risk assessment in prostate cancer screening, and I argue that prostate cancer diagnosis is pervaded by epistemic risks that presuppose value judgments (Biddle and Kukla 2017). These risks include inductive risks and data formation risks—in particular, risks involved in the assignment of Gleason scores to biopsied cells. The pervasiveness of these epistemic risks creates significant and under-explored difficulties for physician-patient communication and, more generally, the achievement of autonomous patient decision making. I will examine the obstacles that must be overcome if genuine respect for patient autonomy is to be reached and argue that the obstacles are sufficiently high as to call into question the feasibility of the updated guidelines.

Biddle, Justin B. and Rebecca Kukla. 2017. “The Geography of Epistemic Risk.” In *Exploring Inductive Risk: Case Studies of Values in Science*, ed. K. Elliott and T. Richards, 215-237. Oxford: Oxford University Press.

United States Preventative Services Task Force (USPSTF). 2016. “Draft Recommendation Statement. Prostate Cancer: Screening.”

<https://www.uspreventiveservicestaskforce.org/Page/Document/RecommendationStatementDraft/prostate-cancer-screening1>.

Welch, H., et al. 2011. *Overdiagnosed: Making People Sick in the Pursuit of Health*. Boston: Beacon.

Dr. Watson: The Impending Automation of Diagnosis and Treatment

Bennett Holman – Yonsei University, South Korea

Last year may be remembered as the pivotal point for artificial “deep learning” and medicine. A large number of different labs have used Artificial intelligence (AI) to augment some portion of medical practice, most notably in diagnosis and prognosis. I will first review the recent accomplishments of deep-learning AI in the medical field, including: the landmark work of Esteva et al. (2017) which showed that AI could learn to diagnose skin cancer better than a dermatologist; extensions of similar projects into detecting breast cancer (Liu et al., 2017); Oakden-Rayner et al.’s (2017) work showing AI could create its own ontological categories for patient risk; and through analyzing tumor DNA

identify more possible sites for intervention (Wrzeszczynski et al., 2017). I will next argue that a foreseeable progression of this technology is to begin automating treatment decisions. Whether this development is positive or negative depends on the specific details of who develops this technology and how it is used. I will not attempt to predict the future, but I will run out some emerging trends to their logical conclusions and identify some possible pitfalls of the gradual elimination of human judgment from medical practice. In particular some problems could become significantly worse. It is the essence of deep learning AI that reasons for its outcomes are opaque. Many researchers have shown that industry has been adept at causing confusion by advancing alternative narratives (e.g. Oreskes and Conway, 2010), but at the very least with traditional research there were assumptions that could, in principle, be assessed. With this deep learning AI there are no such luxuries. On the other hand, I will argue that properly implemented deep learning solves a number of pernicious problems with both the technical and the social hindrances to reliable medical judgments (e.g. the end to a necessary reliance on industry data). Given the multiple possible routes that such technology I argue that consideration of how medical AI should develop is an issue that will not wait and thus demands immediate critical attention of philosophy of science in practice.

Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. "Dermatologist-level classification of skin cancer with deep neural networks." *Nature* 542, no. 7639 (2017): 115-118.

Liu, Y., Gadepalli, K., Norouzi, M., Dahl, G. E., Kohlberger, T., Boyko, A., ... & Hipp, J. D. (2017). Detecting cancer metastases on gigapixel pathology images. *arXiv preprint arXiv:1703.02442*.

Oakden-Rayner, L., Carneiro, G., Bessen, T., Nascimento, J. C., Bradley, A. P., & Palmer, L. J. (2017). Precision Radiology: Predicting longevity using feature engineering and deep learning methods in a radiomics framework. *Scientific Reports*, 7(1), 1648

Oreskes, N. and E. Conway (2010), *Merchants of Doubt*. New York: Bloomsbury Press

Wrzeszczynski, K. O., Frank, M. O., Koyama, T., Rhrissorrakrai, K., Robine, N., Utro, F., ... & Vacic, V. (2017). Comparing sequencing assays and human-machine analyses in actionable genomics for glioblastoma. *Neurology Genetics*, 3(4), e164.

Diagnosis: representing or intervening?

Stephen John – University of Cambridge, United Kingdom

A recurrent theme in recent work on the diagnosis and treatment of cancer is that there is “overdiagnosis”. Unlike “misdiagnosis”, “overdiagnosis” does not refer to straightforward epistemic errors, but, rather, to a systemic tendency to diagnose individuals as requiring treatment, when, in fact, such treatment is medically unnecessary. One obvious way in which to reduce overdiagnosis rates is by changing the diagnostic criteria used to identify lesions and other growths as cancerous. This paper discusses an under-explored topic raised by these debates: how we should conceptualise the act of medical diagnosis. It argues that we should think about this question in terms of speech-act theory, hence suggesting a novel way into broader debates about the relationship between non-epistemic values and scientific justification.

Section 1 uses Hacking’s classic work to outline two ways of thinking about diagnoses: first, as attempts to represent the world; second, as attempts to intervene in the world. I suggest that most discussions of diagnosis start from the first perspective: they assume that a diagnosis is a more-or-less accurate attempt to describe some fact about an individual. Debates around overdiagnosis, however, seem to assume something closer to the second perspective, according to which diagnostic guidelines should be decided according to their non-epistemic effects. I suggest that both perspectives seem incomplete; the former cannot handle debates over overdiagnosis, but the latter cannot explain how we might engage in such activities as retrospective diagnosis.

Section 2 sketches a way of resolving this problem. I set out a model of diagnosis as a particular kind of speech act, an exercitive, which involves changing the status of some individual within the healthcare system. Any specific instance of such a speech act should be governed by sets of epistemic rules for representing phenomena, plus general constraints on permissible speech. However, I also suggest that, at a meta-level, those epistemic rules themselves should be assessed not only in terms of how well they capture the underlying phenomena, but their predictable, systemic effects. I illustrate this model through in-depth discussion of debates over early detection of cancer. Sections 3 and 4 develop the broader implications of my arguments. In Section 3, I suggest that my approach does not reduce to familiar “normativist” theories of disease, but, rather, is compatible with a form of “naturalism”, just so long as we are careful to distinguish the theoretical goals of biology from the practical goals of medical practice. In Section 4, I set out how these arguments relate to debates over the proper role of “non-epistemic values” in

scientific justification, suggesting that we need to distinguish the proper role of values in one-off judgments and the proper role of values in setting our epistemic standards in the first place.

Symposium: Phenomenology and Perspectivism in Science: How Should We Think of Scientific Realism and Scientific Practice?

Organizers: Annamaria Carusi, Franklin Jacoby, Themistoklis Pantazakos, Harald Wiltsche.

Contributors:

Annamaria Carusi.

Franklin Jacoby.

Themistoklis Pantazakos & Harald Wiltsche.

On the one hand, scientific realism traditionally commits us to a positive epistemic attitude towards the truth of theories, the existence of unobservables, and the literal interpretation of scientific claims. On the other hand, the objectivity and mind-independence associated with scientific realism are often in tension with certain elements of science. Scientific theory and practice are human pursuits, and as such contextually situated activities, carried out from a certain perspective. Just what this notion of ‘perspective’ implies, and especially what repercussions it has for science’s claim to the truth has historically been and currently is a matter of great contention within the analytic philosophy of science. This symposium addresses the above issues with a series of talks. Multiple theoretical avenues are utilised, with special focus on scientific practice.

Franklin Jacoby, of the perspectival realism school, argues for a historically situated and contextual notion of perspective, which is analysed in the way scientists evaluate evidence and apply concepts against a theoretical background. This background can vary, but at the same time, and in keeping with a mild form of scientific realism, scientists are argued to display a relatively stable intersubjectivity regarding the norms of evaluating evidence, and to share a common conceptual system. Because of this shared background, perspectives are not incommensurable. These conclusions are displayed through analysis of specific scientific practices.

Harald Wiltsche and Themistoklis Pantazakos tackle the symposium topic from a more continentally-oriented, phenomenological approach, the tools from which they attempt to introduce to the traditionally analytic framework of the scientific realism debate. Wiltsche examines Hermann Weyl’s philosophy of physics, which he utilises to draw conclusions about scientific realism. Building

on a case study, Wiltsche discusses Weyl's notion of perspectivity and its relevance for our philosophical understanding of mathematical models in theoretical physics. With regard to the scientific realism debate, the upshot is that Weyl goes beyond the contemporary dichotomy between realism and anti-realism by introducing a phenomenological notion of constitution.

Pantazakos examines perspective not as an agent's different ways of looking at the same object, but as a shift of conceptual scheme across subjects. He argues against Donald Davidson's contention that there cannot be more than one conceptual schemes. He draws some tentative conclusions regarding scientific realism and explores what this means in practice and how it should inform it via the examination of cases within the autism spectrum (ASD).

Scientific realism is often connected to arguments concerning observation and observables in science, and therefore to scientific vision. Annamaria Carusi argues that perspective and perspectivalism, terms which are rooted in the discourse of vision, do not sufficiently take into account the intra-constitution of vision as seeing and as seen. This intra-constitution is not only a matter of conceptual schemes that may be in operation, but of what Merleau-Ponty talked of as styles. Building on this notion, Carusi presents an argument for styles of realism.

Realisms, Perspectives and Styles

Annamaria Carusi – University of Sheffield, United Kingdom

Scientific realism is often connected to arguments concerning observation and observables in science, and therefore to scientific vision. Even when it is not directly connected to vision, it is metaphorically associated with vision. The 'view for nowhere' is an analogy often used for realism, while talk of perspective brings to the fore that views are never from nowhere, but are made from specific positions, and are always partial, in two senses: incompleteness and bias. Philosophers of science such as Donna Haraway and Sandra Harding have stressed the multiplicity, heterogeneity and interestedness of perspectives that cluster around scientific domains. In this presentation, I analyse the connections between perspective and realism, and argue that even though acknowledging perspectives may appear to be a way of attenuating the 'view from nowhere' account of scientific realism, it remains locked into a discourse of vision that is still predicated upon a distinction between see-er and seen. The presentation outlines a different account of vision, drawn from Maurice Merleau-Ponty's phenomenology, which instead 'gets underneath' perspective itself. Vision, for Merleau-Ponty, implies not just a perspective, but a 'stylised deformation' of the visual field, in which there is a

'complicity' between the act of seeing and whatever is seen. Neither of these features of vision are due to a conceptual scheme, as such, but rather are due to the embodied nature of vision. Merleau-Ponty gives a special priority to vision in certain forms of art to articulate his theory of vision, and it may seem that this has little to do with science. In this presentation I will discuss examples of vision in science in order to show that this is not the case. I set out a theory of vision in science inspired by Merleau-Ponty, that shows the extent to which there is a continuity between 'poetic vision' and scientific vision, and argue that this shows vision and its objects, in science, to be intra-constituted. Finally, I discuss the implications of this theory for scientific realism.

Perspectives, scientific practice, and realism

Franklin Jacoby – University of Edinburgh, United Kingdom

How should we understand disagreement in science? This elements of scientific practice have long provided motivation for non-realist accounts of science, a recent example of which is Chang's normative pluralism (*Is Water H2O?*, 2012). Disagreeing scientists, this view proposes, are often members of incommensurable and coherent systems of practice, which are composed of independent goals, methods, and standards of success. Because they are incommensurable, there is no independent way to arbitrate between systems. Each system can only be evaluated internally by its own standards. We should, Chang argues, be pluralist about these systems and not endorse an imperialistic monism.

Drawing on Giere's (2006) *Scientific Perspectivism*, I suggest another way of thinking about disagreement that has a stronger realist flavour, but with a continued emphasis on science as an activity. Perspectivism is meant to strike a middle ground between realism and relativism. Rather than claim that a mature theory is approximately true, Giere suggests that "According to this well confirmed theory or reliable instrument, the world seems to be roughly such and such" (Giere 2006, p. 6).

I will argue that we should think of a perspective as an evaluation of the evidence that informs scientific concept applications. A perspective is historically-situated and contextual in that when scientists evaluate evidence, they do so on the basis of their background knowledge and theory. The evidence thus evaluated informs how scientists apply concepts to a common, shared world. In this sense perspectives can change and vary between scientists because scientific knowledge changes and different scientists come with different background knowledge. At the same time, and in keeping with a mild form of realism, scientists share general norms of evaluating evidence and

share a common conceptual system. Because of this shared background, perspectives are not incommensurable.

This view achieves three things. Firstly, it avoids the imperialistic “view from nowhere” that Chang and others seek to avoid by situating scientific debates in a historical context that should be evaluated by those historical standards. Instead, this offers a view of science that Massimi calls “from within” (Massimi in Pfeifer and Couch, 2016). A perspective allows us to say, with Giere, that from this historically-situated perspective, the world appears to be such-and-such. Secondly, a perspective offers a more realist flavour by allowing for the rational resolution to scientific debate on the basis of objective, shared norms of evaluating evidence. Science, by this view, can be about a single, objective world without imperialism. Thirdly, perspectivism gives the active elements of science their due by acknowledging that using evidence and applying concepts are actions in a historically-situated practice.

Constitution and Perspectivity: Hermann Weyl’s Phenomenological Non-Realism

Themistoklis Pantazakos – University College London, United Kingdom

Harald Wiltsche – University of Graz, Austria

Scientific realism (SR) is often presented as the conjunction of three sub-theses: first, the metaphysical thesis according to which the world has a definitive and mind-independent structure; second, the semantic thesis according to which theories are truth-valued descriptions of their intended domain, both observable and unobservable; and third, the epistemic thesis according to which science provides us with the means to determine the truth-values of our theoretical descriptions of reality. Looking back at how the debate has evolved over the last decades, philosophical attention has been paid almost exclusively to the second and third sub-thesis. While earlier discussions were primarily concerned with the possibility of expressing theories in a purely observational vocabulary, the debate took a distinctively epistemic turn since the 1980ies. What is also noteworthy is that the discussion between SR and scientific anti-realism (SAR) seems to have reached a stalemate. As argumentative strategies on both sides have become increasingly sophisticated, the prospects of breaking the deadlock in a non-circular fashion do not appear very bright. This has led some philosophers to dismiss the entire issue as artificial.

I accept the reality of a stalemate between contemporary versions of SR and SAR. On my view, however, its existence is not an indication that the issue is unworthy of philosophical effort to begin with. Rather, as I shall argue, the

stalemate results from the lack of attention that has been paid to the metaphysical sub-thesis, an assumption that is shared both by proponents of SR and SAR. In order to make my point more vivid, I will discuss the work of Hermann Weyl, one of the premier mathematicians and theoretical physicists of the 20th century who also invested much energy into reflecting on his scientific practice from a philosophical point of view. Building on a concrete case study, I shall show how Weyl uses the phenomenological notions of constitution and perspectivity in order to develop an account that is based on a more qualified understanding of the metaphysical thesis. By examining the nature of scientific cognition from within a phenomenological framework, Weyl comes to an understanding of science and scientific practice that goes beyond the dichotomy between SR and SAR.

Symposium: Coherence as an Indicator of Quality in Scientific Practices

Organizer: Hasok Chang.

Contributors:

Hasok Chang.

Lena Soler.

Sjoerd Zwart.

Among philosophers who seek a full understanding of scientific practices, two things are commonly accepted: knowledge is not merely propositional, and there is no rock-bottom foundation of empirical knowledge. But it is not clear how the quality of non-propositional and non-foundationalist knowledge should be evaluated. This is in contrast to the illusory comfort of the traditional view, with truth-as-correspondence as a straightforward but inoperable epistemic ideal.

Coherence has been proposed as a plausible epistemic ideal by some drivers of the “practice turn” in the philosophy of science (e.g. Hacking 1992, Pickering 1995). Traditionally, with interest primarily directed toward propositions, coherence referred to logical consistency or some broader mutual support-relation among a set of propositions. Such a conception is clearly too narrow for the evaluation of practices. With the practice turn, targets of interest include not only propositions but experiments, material devices, multiple types of scientific activities, paradigms, and socio-technical systems. The epistemic quality of these targets tend to be understood in more concretely applicable terms, such as reliability, success, robustness, efficiency, etc. In this context, it is much more difficult to specify what coherence means, how it works, and

what its exact relation is to reliability, success, or the like. This symposium seeks to contribute to the development of feasible precise notions of coherence.

If we say that coherence is not a relation between propositions, we could say that it is a relationship among instances of another kind of thing; this might be called the “homogeneous view”. Alternatively, one can advance a “heterogeneous view”, seeing coherence as a relation holding among diverse types of elements of scientific practice. Some would argue that both of these views are too reductionistic and limiting; instead, they want to see coherence as a property holistically possessed by a system as a whole.

The proposed symposium will consist of three papers. Léna Soler will examine the relation between coherence and reliability, focusing on the particular case of the reliability of an experimental proof. In so doing, she will develop her ideas on how coherence consists in a “symbiosis” (in Pickering’s words) between multiple types of elements of scientific practice, and how the corresponding “enlarged coherence” (as Hacking terms it) involves diverse kinds of “glue” to bind the elements together. She will illustrate these points through the case of the discovery of weak neutral currents in particle physics. Hasok Chang will further develop his previous proposal that coherence should be seen as a relation among actions (Chang 2014): a certain harmoniousness holding among the operations that together constitute an activity, which is conducive to the successful achievement of the aims of that activity. Spelling out this notion will also require close attention to the ontology of action, ability and knowledge. A dilemma remains: if we define coherence strictly in terms of success, it becomes a redundant concept; but otherwise is it an ill-defined and unobservable quality?

Sjoerd Zwart will offer a contrasting view, according to which Chang’s and Soler’s analyses come out looking too reductionistic. He considers coherence as an emergent property of a whole system. He argues that the relations that do exist between the elements are captured better by notions other than coherence, and that an emergent concept of coherence can do indispensable work in many practical and theoretical areas. His points will be illustrated by particular attention to the coherence of engineering projects.

Chang, Hasok. 2014. “Epistemic Activities and Systems of Practice: Units of Analysis in Philosophy of Science After the Practice Turn.” In Léna Soler, Sjoerd Zwart, Michael Lynch and Vincent Israel-Jost, eds., *Science After the Practice Turn in the Philosophy, History and Social Studies of Science*, 67–79. London and Abingdon: Routledge.

Hacking, Ian. 1992. “The Self-Vindication of the Laboratory Sciences.” In Andrew Pickering, ed., *Science as Practice and Culture*, 29–64. Chicago: University of Chicago Press.

Pickering, Andrew. 1995. *The Mangle of Practice: Time, Agency, and Science*. Chicago: University of Chicago Press.

Operational Coherence as the Effective Harmony of Actions

Hasok Chang – University of Cambridge, United Kingdom

In some recent works (Chang 2017; 2018) I have proposed to ground pragmatist notions of truth and reality in the concept of “operational coherence”. Whether or not such proposals will work out, much work is needed to develop more fully and sharply this notion of coherence, which I had defined as follows: an activity is operationally coherent if and only if there is a harmonious relationship among the operations that constitute the activity, which is conducive to the successful achievement of the aims of that activity. (Similarly, on a larger scale, coherence is harmoniousness among the activities that constitute a system of practice.)

The key insight I want to preserve and develop, building also on other previous work (e.g., Chang 2014), is that coherence should be seen as a relationship among actions, not among propositions or entities. Operational coherence consists in various actions coming together in an effective way towards the achievement of one’s aims. It comes in degrees and different shapes, and it is necessarily a less precise concept than consistency, which comes well-defined through logical axioms. An important part of my proposal is to keep in mind the aims that scientists are trying to achieve in each and every situation. The presence of an identifiable aim distinguishes actions from mere physical happenings involving human bodies, and it also places knowledge firmly in the realm of actions.

There are two main issues I want to address further. First, what exactly is the relationship between coherence and success? As I see coherence as conducive to success, I want to say that the concrete realization of a coherent activity is successful, *ceteris paribus*. And this latter condition serves as an indirect criterion for the judgement of coherence. But why make that link indirect, rather than simply identifying coherence with success? I want to allow that an activity may be successful by accident, rather than by its proper coherence; conversely, a perfectly coherent activity may fail because of some accidental circumstances. But this move threatens to make coherence an ill-defined and unobservable quality, and I will need to consider the philosophical consequences of that. There is a dilemma: if coherence is defined strictly in terms of success, then it is a redundant concept; if it is not, then it is inaccessible.

The second remaining issue is the very ontology of activities. In my previous

work I conceived a hierarchical nesting of doings: operations making up an activity, and activities making up a system of practice (Chang 2014). But there is another way of conceiving an activity, which is to see it as composed of a heterogeneous set of elements, including the agents' beliefs, abilities, bodily movements, social rules, and material objects. This is surely a sensible perspective, but how exactly do we conceive operational coherence in such framing of action? I will draw on some of the recent literature in the philosophy of action that stress the role of the active agent (e.g. Hornsby 2007), and also some rather neglected classics such as Gilbert Ryle's (1945/46) arguments concerning "knowing how".

Chang, Hasok. 2014. "Epistemic Activities and Systems of Practice: Units of Analysis in Philosophy of Science After the Practice Turn", in L. Soler, S. Zwart, M. Lynch and V. Israel-Jost, eds., *Science After the Practice Turn in the Philosophy, History and Social Studies of Science* (London and Abingdon: Routledge), 67–79.

Chang, Hasok. 2017. "Operational Coherence as the Source of Truth", *Proceedings of the Aristotelian Society* 117: 103–122.

Chang, Hasok. 2018. "Is Pluralism Compatible with Scientific Realism?", in Juha Saatsi, ed., *The Routledge Handbook of Scientific Realism* (London: Routledge), 176–186.

Hornsby, Jennifer. 2007: 'Knowledge and Abilities in Action', in C. Kanzian and E. Runggaldier (eds.), *Cultures. Conflict - Analysis – Dialogue* (Frankfurt: Ontos Verlag), 165–180.

Ryle, Gilbert. 1945/46. "Knowing How and Knowing That: The Presidential Address", *Proceedings of the Aristotelian Society*, new series, 46: 1–16.

Philosophy. <https://plato.stanford.edu/entries/molecular-genetics/>

What reliability judgments about experimental proofs owe to coherence

Léna Soler – Université de Lorraine, France

A pervasive intuition, at work in many different philosophical approaches of science has been that coherence is the main, if not the only criterion available to assess the "quality" of any "ingredient" X of science. However, the intuition proved difficult to turn into a precise characterization, especially within the practice turn, once it is recognized that coherence cannot be reduced to the absence of contradiction or some other relation between propositions. The talk focuses on the case in which X is an experimental proof, relying on the example of the experimental detection of weak neutral currents (NC) in the

mid-1970 (Soler 2012). It aims to discuss what reliability judgments about such kind of X owes to coherence understood in an enlarged sense, as a “symbiosis” between multiple elements of scientific practice (Hacking 1992, Pickering 1995). The issue of how to conceptualize the nature of the elements and the kinds of “glue” between them is part of the difficulty of characterizing the working of coherence in science, and will be addressed in the talk. First, a common presentation of the “discovery of NC” in the mid-1970 is considered, according to which the experimental result R = “existence of NC” can be taken as reliable since three different experimental proofs – or as I prefer to say for reasons that will be explained, three experimental “derivations” $D1$, $D2$ and $D3$ – independently support R (for example, $D1$ derives R from an experiment using a bubble chamber as detector, whereas $D2$ derives R from an experiment using electronic detectors). According to this presentation, the historical stabilized situation is conceptualized, at a certain scale, through a relatively simple four-ingredients scheme – called a “robustness scheme” in honor to Wimsatt’s definition of robustness – in which $D1$, $D2$ and $D3$ converge on one and the same R . Starting from this conceptualization, I discuss the working of such a robustness scheme, and analyze in what sense, and to what extent, the reliability judgments about each of the four ingredients emerge from the circumstance that all of them are embedded in a coherent holistic unit (or scientific symbiosis).

Second, I focus on one of the experimental derivations, $D1$, and ask the question of its reliability. As one ingredient of the previous robustness scheme, $D1$ owes (at least part of) its strength from the “external” support provided by the other ingredients involved in the scheme. But does this “external” source of strength exhaust what we take into account when we assess the reliability of an experimental proof like $D1$? Intuitively, we feel that we should consider something like the “internal” strength of $D1$ “in itself”, i.e., independently of the circumstance that historically, $D1$ has been embedded in a robustness scheme. To scrutinize what lurks behind this intuition about “internal reliability”, I open the black box of the experimental derivation $D1$ and offer a (partial and schematic) conceptualization of what we find when we look “inside”. I show that what we find inside can be analyzed, once again, as a holistic coherent unit or symbiosis, but in this case, we have an architecture of a much more complex type than the robustness scheme. Some insights are then provided about the nature of the global good fit here involved: about the way it emerges from multiple more local sub-fits, and about the kinds of local symbiotic relations that constitute it.

Finally, I explain in what sense the previous characterization, although already very complex, is a highly simplified conceptualization of the targeted reality. Insights are provided about what should be added to the picture, which point to essential difficulties regarding the question “the coherence of what?”.

Hacking, Ian. 1992. "The Self-Vindication of the Laboratory Sciences." In Andrew Pickering, ed., *Science as Practice and Culture*, 29–64. Chicago: University of Chicago Press.

Pickering, Andrew. 1995. *The Mangle of Practice: Time, Agency, and Science*. Chicago: University of Chicago Press.

Soler, Léna. 2012. "Robustness of Results and Robustness of Derivations: The Internal Architecture of a Solid Experimental Proof." In Léna Soler, Emiliano Trizio, Thomas Nickles & William C. Wimsatt, eds., *Characterizing the Robustness of Science after the Practice Turn in Philosophy of Science*, Dordrecht: Springer, Boston Studies in the Philosophy of Science 392, 227-266.

Coherence as Emergent Property of Large Ensembles

Sjoerd Zwart – Delft University of Technology, Netherlands

The main claim defended and illustrated in this paper is that coherence in scientific and engineering practices is most sensibly attributed only to relatively large compound ensembles of elements some of which should be human actions. Thus, I oppose attempts to characterize coherence of an ensemble in terms of the coherence between its individual elements—such as atomic actions of the type "lighting a match" or "hitting a nail on the head". Like temperature, coherence only becomes apparent on a larger scale; it is primarily an emergent property of compound ensembles that cover at least a number of human actions with distinct objectives.

Reasons for this point of view are threefold. First, for the assessment of the relationship between the atomic elements of the ensemble the coherence notion is too vague; it resists consistently endeavors to give it a clear and convincing explanation on the level of atomic elements (Thagard, 2002; Bratman, 1999; Chang, 2017). Second, most coherence associated relationships between elements like actions, means, ends, propositions, processes, 'facts', are more forcefully and exactly explained in other terms than coherence. Practitioners assess these relationships using terms like efficiency, effectiveness, verification, validation, cause, effect, probabilities, explanation, etc. Renaming these notions in terms of coherence is not only an innocent play of words; it also runs the risk of introducing possibly harmful ambiguities. Third, coherence on the level of ensembles need not be explained in terms of coherence between its elements. Taking coherence as an emergent property does not impair its indispensable work in many practical and theoretical areas in philosophy, science and engineering (such as in Kuhn's paradigms; coherence theory of truth; Rawls' reflective equilibrium).

When we accept the central claim of this paper, the main work is still left to be done. We still have to show how coherence on the aggregate level is to be explained by other types of relationships between the elements of on the basic level. Part of the practices of science and engineering studies then becomes the exploration of the relationships between the elements that aggregate into the (in)coherence of the overarching ensemble. To do so, I will turn to empirical studies in engineering research. One concerns the introduction, research and development of Nereda[®], a cutting edge wastewater treatment technology (Kreuk et al 2010); the other focuses on engineering PhD research projects. In the first case coherence emerges from using a (series of) experiment(s) for different purposes (a scientific and an engineering one). And regarding the second coherence is strongly correlated with the hierarchical design of the research plans, which are built up out of six types of atomic engineering projects (Zwart and de Vries 2016). For instance, a four-year research project eventually aiming at engineering know-how may consist of various types of subprojects: a study of the relevant scientific and engineering theories; research into the pertinent context and current practices; conceiving possible working principles and comparing them; defining proposals for intervention; testing the proposed intervention in a laboratory and in practice, including possible collateral damage. Such a project is coherent if all subprojects explicitly support the ultimate goal and are mutually supportive.

Bratman, M., & Stanford University Press. (1999). *Intention, plans, and practical reason*. Stanford: Stanford University.

Chang, H. (2017). *Operational Coherence as the Source of Truth*. In *Proceedings of the Aristotelian Society* (Vol. 117, pp. 103–122). Oxford University Press.

Kreuk, M de, Poel, I.R van de, Zwart, S.D. and M.C.M. van Loosdrecht (2010) 'Ethics in innovation: cooperation and tension.' In I.R. van de Poel and D.E. Goldberg (Eds.), *Philosophy and engineering: an emerging agenda*. Dordrecht / New York / Berlin: Springer, pp. 215-226.

Thagard, P. (2002). *Coherence in Thought and Action*. MIT Press.

Zwart, S.D. and M. de Vries (2016) 'Methodological Classification of Innovative Engineering Projects.' In: Franssen, M., P. E. Vermaas, P. Kroes and A Meijers (Eds.), *Philosophy of Technology after the Empirical Turn*. Dordrecht / New York / Berlin: Springer.

Symposium: Assessing Economic Expertise

Organizers: François Claveau, Julian Reiss.

Contributors:

François Claveau & Jérémie Dion.

Carlos Martini & Anita Välikangas.

Julian Reiss.

Economic experts are thriving. They seem every year more present in the media. And this is only the visible part of the iceberg: economists permeate modern bureaucracies by writing reports, giving counsel, and often pulling themselves the policy levers. They thrive even though they are scorned at. Often depicted as contemporary ideologues dressed up as scientists, they persist unperturbed.

Experts are socially recognized specialists of a domain that are given some authority to affect public policy. Since one main way for scientists to impact society is to play this expert role, the philosophy of science in practice must partly be the philosophy of expertise. In recent years, philosophers of science have indeed turned their attention to the assessment of scientific expertise. Since economics is both a massive supplier of expertise and a scientific discipline with a stained track record, the assessment of economic expertise should be among the priorities of philosophers of science in practice. This symposium is a contribution to this assessment. Its general question is: Under which conditions should laypersons can reasonably trust economic experts? Julian Reiss replies to recent contributions that argue that the authority of experts in society should be strengthened. He responds that the conditions for a sensible deference to experts are not present in the case of economics. The other contributors focus on specific organizations of economic expertise by combining conceptual and empirical analyses. Carlo Martini and Anita Välikangas study markers of trustworthiness in think tanks. Emmanuel Carré assesses central banks in their role of financial supervisors: Can we trust them to keep financial exuberance in check? Finally, François Claveau and Jérémie Dion also assess central banks and find tensions between two types of expertise: testimonial and regulatory.

Regulatory versus testimonial expertise: the case of central banking

François Claveau – Université de Sherbrooke, Canada

Jérémie Dion – Université de Sherbrooke, Canada

Central banks are expert communities who are supposed to “promote the good of the people” (according to the Bank of England for example). The “people” accordingly rely on central banks to perform key societal functions. The reliance of laypersons is rational only if these institutions are trustworthy. But how is it possible for laypersons to form justified beliefs in the trustworthiness of central banks? This question raises a well-known issue in the lay-expert relationship (Hardwig 1991). By addressing the issue in the case of central banking expertise, this paper contributes both to our general understanding of rational trust in experts and to a timely epistemic assessment of institutions which have significantly grown in influence since the 2007-08 crisis.

Our contribution centers on the distinction between having regulatory expertise and having testimonial expertise in a domain. As applied to the domain of central banking, members of central banks act as regulatory experts when they make decisions on, most importantly, which monetary policy to pursue. They share this type of expertise with other public regulators such as members of drug authorization agencies. Yet, members of central banks also play a prominent role as testimonial experts, that is as providers of information to laypersons on how central banking works and how it should work.

Most testimonial experts are not regulatory experts – e.g., an academic researcher working on cross-country differences in drug authorization procedures is typically at arm’s length with the actual regulators. In addition, most regulatory communities do produce some testimonial expertise on their domain, but they do not spend a large fraction of their resources on this task and are not the main source of information on their domain. Central banks stand out in this respect: in the last decades, they have become dominant testimonial experts on central banking through the “scientization” of their identity (Marcussen 2009; Mudge-Vauchez 2016). For instance, more than half of the articles in the three main academic journals on central banking are now authored by staff from central banks (Claveau, Dietsch and Fontan, forthcoming).

The main goal of this paper is to assess how, for such a community, the imperatives of the two types of expertise partly reinforce and partly undermine each other. We use both in-depth empirical research (institutional analysis and computer-assisted textual analysis) and conceptual analysis to pursue this goal.

Our main thesis is twofold. On the one hand, the explicit goal of central banks in ramping up internal research – i.e., becoming better regulatory experts – is likely to be met. Indeed, the standard tools (e.g. Longino 1990; Goldman 2001) for laypersons to assess the trustworthiness of experts pronounce in this direction. On the other hand, the scientization of central banks and its accompanying research concentration have worrisome consequences for testimonial expertise on some central banking topic. Most importantly, pronouncements on the proper delegation of powers to central banks – a topic that is trending up in public discussions – suffer from a serious credibility deficit when they come from central bankers themselves (because of concerns about conflict of interest).

Beyond the empirical specifics of the central banking community, this paper is meant as a contribution to the reflection on the rational trust of laypersons in experts. It advances this reflection by analyzing how rational trust should be differentially affected by the intermingling of regulatory and testimonial expertise.

Claveau, François, Dietsch, Peter and Clément Fontan (forthcoming) *The Ethics of Central Banking*. In *Handbook in Ethics and Public Policy*.

Goldman, A. I. (2001). *Experts: Which Ones Should You Trust?* *Philosophy and Phenomenological Research*, 63(1), 85–110.

Hardwig, J. (1991). *The Role of Trust in Knowledge*. *The Journal of Philosophy*, 88(12), 693-708.

Longino, H. E. (1990). *Science as Social Knowledge*. Princeton: Princeton University Press.

Marcussen, M. (2009). *Scientization of Central Banking: The Politics of A-Politicization*. In K. Dyson and M. Marcussen (eds.), *Central Banks in the Age of the Euro: Europeanization, Convergence, and Power* (p. 373-390). Oxford: Oxford University Press.

Mudge, S. L. and Vauchez, A. (2016). *Fielding supranationalism: the European Central Bank as a field effect*. *The Sociological Review Monographs*, 64(2), 146-169.

Expertise and Trust in Think-Tank Research

Carlo Martini – UniSR, Italy

Anita Välikangas – University of Helsinki, Finland

In this article we investigate how transparency works as a vehicle of trust in think tank research. Trust is highly valued in science, because both scientists and science-users tend to benefit when society trusts science to deliver the

fruits it promises. Scientists go to great lengths to gain the public's trust, including fostering a culture of transparent expertise, where scientists' public professional profiles are guarantors of the trustworthiness of their professional activities. In this paper we focus on an important repository of scientific research, that is, think tanks. In a number of key fields of research, like economics and health science, think tanks produce extremely policy-relevant knowledge, hence the need to study the culture of expertise and transparency in think tank research. Our research provides an empirical survey of the practices of transparency in communicating expertise in think tanks.

It has been shown repeatedly that trust is fundamental to the development and progress of research itself, that is, to knowledge creation. Scientists need to trust one another to push their research forward (Hardwig, 1985, 1991), and modern science needs public trust both for its material and moral sustainment. Indeed, the mechanisms through which scientists communicate and cultivate a relation of trust within and without their profession has been researched extensively in monographies (see Gross, 2002; Russell, 2010) and is a topic of continuous research in major journals like *Science Communication* (Sage) and *Public Understanding of Science* (Sage).

Think tanks are one of the major sources of non-academic research in several policy-oriented research fields, from economics to public health. Thanks to their focus on communication and policy, think tanks are typically positioned very close to the political arena, and thus more effective at influencing policy decisions. While social research sometimes does not classify think tanks as research bodies, but rather as intermediary organizations between academic and policymaking communities, in this paper we wish to focus our case study on think tanks that maintain research-related activities. So far, the emphasis of research has been on how think tanks influence policy and society, instead of how they maintain the quality of their scientific research (see e.g., Abelson, 2002; Weidenbaum, 2010).

The empirical study we provide in this paper surveys a sample of 293 individual authors of think tank reports listed as original research. We are interested in understanding how transparent the expertise behind the reports is, and whether think tanks promote public trust in their research by furthering what we call "transparency of expertise". We provide a classification of what kind of information is available for the authors we select and we look for explicit information about the culture of transparency of each think tank we survey.

Against Epistocracy

Julian Reiss – Durham University, United Kingdom

Is there such a thing as an ‘economics expert’? While much of the earlier literature on the role of experts in society has focused on limiting the power of experts by subjecting it to democratic control (most prominently, perhaps, in Paul Feyerabend’s *Science in a Democratic Society*), a number of more recent contributions argue in favour of something that comes close to the exact opposite: the subjugation of democracy to scientific control, and control by economists and other social scientists in particular. In this paper I focus on the two books *Why Democracies Need Science* by Harry Collins and Robert Evans and *Against Democracy* by Jason Brennan, both of which advocate the creation of new, science-strengthening institutions: the former, a committee of ‘owls’ — scientific experts who assess and certify the quality of a scientific consensus of some policy-relevant matter; the latter, the replacement of the ‘one person – one vote’ principle by a principle according to which a person’s voting rights are, in part, made dependent on the person’s expertise in scientific (especially, social-scientific and economic) matters. Against these, I argue that both kinds of institutions would lead to extremely harmful consequences and urge philosophers to return to the values defended in the earlier literature on experts in society.

One of the major premisses in my argument is a denial of the existence of uncontroversial knowledge in economics and other social sciences. It is because there is no such uncontroversial knowledge that economists and other social scientists cannot be said to have superior judgement in matters of potential social or political relevance such as whether free trade is good for a nation, minimum wages are harmful or unlimited immigration will raise national product. Moreover, even if there appears to be consensus on such a matter, this is highly likely either to be a mirage and appear only because one does not look far enough or in the right places or to have arisen for the wrong reason such as conformism, acceptance of a false theory or of bad social values. Consensus is therefore neither an indicator of truth nor a guide for action.

In sum, while I don’t advocate anything close to a ‘silencing of the experts’, I am very sceptical of recent arguments to the effect that experts’ position in society should be greatly strengthened and instead recommend a close scrutiny of experts’ opinions by democratic processes.

Symposium: Author Meets Critics: Adrian Currie, "Rock, Bone, and Ruin: An Optimist's Guide to the Historical Sciences" (MIT Press, 2018)

Author: Adrian Currie – University of Cambridge, United Kingdom

Commentators:

Leonard Finkelman – Linfield College, United States

Craig Fox – Rotman Institute of Philosophy, Canada

Alison Wylie – University of British Columbia, Canada

In his recent book, *Rock, Bone, and Ruin: An Optimist's Guide to the Historical Sciences* (2018), Adrian Currie develops a systematic view of the practice of the historical sciences. At the center of Currie's picture is the claim that historical scientists are methodologically omnivorous and opportunistic. Much of his work focuses on documenting the variety of different modeling practices and inferential strategies, such as the comparative method and the "exquisite corpse" method, that researchers use to reconstruct prehistory. One theme of Currie's work is that scientists often use background knowledge to provide epistemic scaffolding. With the right background knowledge, one can extract quite a bit of information from even a single fragmentary fossil. Currie's view also connects with topics of current discussion among philosophers of science, such as the nature and role of historical narrative explanation. And some aspects of Currie's picture of historical science will no doubt be controversial. For example, he offers a sustained defense of the epistemic value of speculation in historical science, a line of argument that is likely to meet with some resistance from philosophers and scientists who place a premium on epistemic caution. Currie's book is likely to shape the near-term agenda for philosophical reflection on the practice of the historical sciences.

Currie will begin the session with a 15-minute synopsis of *Rock, Bone, and Ruin*. Each of the three critics will take 10-12 minutes to offer commentary on one aspect of Currie's project. Currie will then take another 10 minutes to reply to the critics, leaving 20-30 minutes for discussion. (This assumes a 90-minute session).

Derek Turner will introduce the session and moderate the discussion. The critics will approach the book from three different directions, with different disciplinary emphases. Alison Wylie brings expertise in the philosophy of archaeology. Leonard Finkelman is a philosopher of biology who is currently pursuing graduate work in paleontology. Craig Fox is a philosopher of science working on the epistemic role that narratives play in historical reconstruction.

Symposium: Philosophy of Astrophysical Practice

Organizer: Jaco de Swart.

Contributors:

Siska Baerdemaeker.

Nora Boyd.

Mauricio Suárez.

Jaco de Swart.

In post-Kuhnian philosophy of science, reflections on the relation between scientific theories and the world made way for laboratory studies, philosophies of experimentation and analyses of scientific practice. In this movement, the largest 'laboratory' in science has received comparatively little attention: the universe. Cosmologists and astrophysicists work with this peculiar laboratory, and their methods and practices aim at producing knowledge of the distant stars, galaxies, clusters, and even the whole universe. Exploring the cosmic laboratory comes with challenges though. The objects and processes of interest are remote in space and time and the signals can be weak and rare, making astrophysical work epistemically far from trivial.

Since astrophysics and cosmology are relatively young empirical sciences, understanding what it means to do a science of the cosmos has been an evolving question throughout the 20th century. Because of this, numerous recent examples are available that illustrate how theories form, models are built, argumentation takes place, instruments are used, and objects of knowledge travel. Potential examples include the rivalry between steady state cosmology and the big bang theory, establishing the extragalactic distance scale, the introduction of modified gravity as an alternative to dark matter, and the measurement of the accelerated expansion of the universe, but the list is much more extensive. With its many controversies and methodological considerations, the astrophysical sciences form a remarkable window into understanding scientific practice and, in particular, the epistemic status of models, observations and evidence.

The philosophy of cosmology and astrophysics is currently experiencing a revival, in parallel with recent discussions on the scientific status of modern cosmological theories (e.g. Smeenk, 2013; Ellis & Silk, 2014; Chamcham et al., 2017). However, a practice-based approach has not yet been well-represented in this revival. If we want to explore the assumptions and methods underlying the astrophysical sciences, it is essential not only to explore theories and results, but also the processes by which these conclusions came to be. The purpose of this session is, indeed, to address the norms and practices of

astrophysics and cosmology philosophically. The contributions in this session span a range of approaches, from historical to analytical, displaying the scope and potential of pursuing a philosophy of astrophysical practice. The central focus is the diverse ways in which 'evidence' is used in this domain – from stellar astrophysics to cosmological model building.

By illustrating the exciting contributions that studies of astrophysics and cosmology can make to the philosophy of science in practice, our wish is to stimulate further work on this rich subject and open broader discussions on its philosophical relevance and connection to practices of other disciplines.

Chamcham, K., Silk, J., Barrow, J. D., & Saunders, S. (Eds.). (2017). *The Philosophy of Cosmology*. Cambridge: Cambridge University Press.

Ellis, G., & Silk, J. (2014). Scientific method: Defend the integrity of physics. *Nature*, 516(7531), pp. 321-3.

Smeenk, C. (2013). *Philosophy of Cosmology*. In R. Batterman (Ed.), *Oxford Handbook for the Philosophy of Physics* (pp. 607-652). Oxford: Oxford University Press.

Integrating evidence in cosmology: the search for dark matter

Siska De Baerdemaeker – University of Pittsburgh, United States

Cosmology at its core is an integrative science. In order to answer one of its central research problems, 'how did the universe evolve from a hot dense state, to the universe that is observed today?', cosmologists draw on a variety of theories, (partial) models and hypotheses, as well as methods and sources of evidence. The process of explanatory, methodological and evidential integration has led to the so-called standard model of cosmology, Λ CDM. This paper explores how cosmologists have overcome difficulties in integrating different sources of evidence, and draws lesson for the future.

I introduce a distinction between two types of evidence in an integrative context: mediated and unmediated evidence. Crudely stated, unmediated evidence originates from the target system itself; its source is the system whose behavior or evolution constitutes the complex research problem. For mediated evidence, the source comes from a different domain than that of the target system, and its applicability to the integrative context therefore needs additional justification. I argue that unmediated evidence should take priority, because of two arguments: the argument from reliability, and the argument from heuristics.

I then apply these arguments to a central cosmological research area where the use of different sources of evidence becomes apparent: dark matter.

Although already proposed in the 1930s, dark matter fully rose to fame in the late 1970s as an explanation for observations of anomalous galaxy rotation curves. Later observations of the Bullet Cluster, as well as accounts of large-scale structure formation added to the evidential support for dark matter, in favor of alternatives like Modified Newtonian Dynamics (MOND). Today, it is one of the main components of Λ CDM.

Although Λ CDM is phenomenologically highly successful, cosmologists currently lack understanding of the fundamental nature of dark matter. Particle physicists have proposed different extensions of the standard model of particle physics in response, most notably Weakly Interacting Massive Particles (WIMPs). In search of WIMPs, several so-called direct detection and production experiments have been set up. These experiments essentially all apply methods from particle physics to the search for WIMPs. So far, none of these experiments have turned out a positive result. With the next generation of experiments becoming even more sensitive, the clock is ticking on the WIMP hypothesis – if anything because the neutrino floor constrains the region where WIMP- and neutrino-interactions can be distinguished in direct detection experiments.

Even if WIMP searches fail to turn out a positive result, I submit that dark matter is here to stay – at least insofar as ‘dark matter’ refers to the explanation of the aforementioned astrophysical phenomena. I defend this view by applying the priority of unmediated evidence from astrophysical observations, over mediated evidence from direct detection searches.

I end with a discussion of how the difference between mediated and unmediated evidence can be applied to other debates in integrative scientific contexts, most notably the history of the cosmological constant Λ and the debate between Maxwell and Kelvin on the age of the earth.

Repurposing Historical Astronomical Data

Nora Boyd – University of Pittsburgh, United States

Historical astronomical records can be valuable, sometimes irreplaceable, for certain research questions. For instance, some astronomical events, such as nearby supernovae, are rare enough that few occurrences have been witnessed since the advent of the telescope. Other astronomical phenomena change only subtly over centuries. In order to use historical records of such phenomena, researchers have implemented clever strategies for coaxing them into epistemic contact with contemporary theory. Having the capacity to use empirical results in contexts besides those that generated them is thus critically important for studying these sorts of astronomical phenomena.

I argue that, in general, in order for some empirical result to serve as a constraint on theorizing in some epistemic context, it must be “well-adapted” to the context of constraint. I defend a precise characterization of well-adaptedness and articulate one strategy by which an empirical result can be repurposed in a new context—using data records and their provenance metadata as the basis for transforming the empirical results codified in those records into useful empirical constraints in the contexts of interest. I develop the notion of “evidential forensics” to capture the clever chains of inference that researchers employ to render some historical records epistemically useful in the present. In particular I discuss three aspects of evidential forensics: assessment of relevance, translation/transformation of information, and circumstantial reasoning.

I present a virtuoso example of evidential forensics in action—the successful transformation of a Babylonian eclipse record from 694 BC into a useful constraint on the evolution of the length of the Earth’s day. The length of a day on Earth has, it turns out, been slowing down. Babylonian eclipse records have helped to put empirical pressure on the idea that the slowdown could be due entirely to tidal braking from the Earth’s gravitational interaction with the Moon. Indeed, fully accounting for the slowdown appears to require contributions from other geophysical processes such as post-glacial rebound and core-mantle coupling. Learning this has required the transformation of results recorded in cuneiform script on broken clay tablets conveying observations recording in utterly defunct spatial temporal units.

In the case I present, certain desiderata that researchers have identified as requisite for an historical record to be useful as a constraint in this context were met. To take just one example, to effectively use the timing information encoded in an eclipse record it must be possible to determine the geographical location from which the observation was made. In addition, the content of the records themselves were deciphered and processed into well-adapted results, and researchers recruited background knowledge about the historical and cultural context in which the observations were originally made in order to make a plausible argument about the timing of the eclipse. I explore the parallels between the epistemology of this sort of evidential forensics with strategies in historiography and archaeology arguing that it shares some features characteristic of each, the most important being that the epistemic utility of the artifact depends crucially on the accessibility of details regarding its provenance.

Evidence for Fictions and Idealizations in Stellar Astrophysics

Mauricio Suárez – Complutense University of Spain, Spain

How do fictions and idealizations differ? An apt characterization has it that while fictions are unconcerned with truth (or faithful representation), idealizations aim at truth in some roundabout way. I argue that they also possess very different functional connections to experimental evidence. Roughly, experimental evidence for idealizations renders them “controllable” when corrections can be brought to bear in order to get those idealizations closer to the true description of their real targets – a procedure known as “de-idealization” in the literature. Thus the evidence for idealized models is “controllable”: It is fine-grained, and modulated. One can accept the evidence for the model without having to accept it as evidence for the unrealistic idealizations within the model. By contrast, evidence for fictions is never “controllable”, since there are no corrections that can be brought to bear in order to get the fiction closer to the target. Rather, a fictional assumption in a model functions merely as an efficient inferential shortcut. That is, a fictional model shows itself valuable in the economy and elegance of the inferences that it promotes to observable quantities or properties of the targets. The experimental evidence provided by the phenomena is coarse-grained, and unmodulated. In other words, accepting “non-controllable evidence” for a model requires accepting it for the model as a whole – including all its fictional assumptions.

I illustrate these theses by means of models of stellar structure in astrophysics. These models typically provide us with four equations (hydrostatic equilibrium, continuity, radiative transfer, and thermal equilibrium) from which the three observable quantities of stellar astrophysics may be derived (these “observable” quantities are the star’s surface temperature, its luminosity, and the spectral distribution of its radiated light). However, the models incorporate a number of assumptions in the derivation of the four equations that are hardly realistic. A star is generally assumed to be any body of gas uniformly constituted by a mixture of hydrogen and helium, bound together by self-gravity, which radiates energy from some internal source. Yet, the models in addition assume: i) uniform chemical composition throughout the star; ii) a spherically symmetrical shape; iii) isolation of the star’s gas from the surrounding interstellar gas; and iv) thermal equilibrium. The assumption of uniform chemical composition (70% hydrogen to 30% helium) can be corrected in different ways, and may be thought of as a controllable idealization. However, the other three assumptions are not controllable. Spherical symmetry and isolation are entirely fictional assumptions which are conducive

to expedient calculation (since together they generate the result that the star's mass is bounded, and determine the layered structure of the star). And I shall focus on the fourth assumption, in particular, which takes the temperature of the gas and the radiation to be identical. I shall argue that evidence can be brought to bear differentially on this assumption, depending on background theory, but that it remains a substantially fictional kind of evidence justified by inferential use rather than fine-grained or modulated.

Making Dark Matter Matter; Anomaly Formation in 1970s Astrophysics

Jaco de Swart – University of Amsterdam, Netherlands

In 1974, two landmark papers were published by independent research groups in the U.S. and Estonia, that concluded on the existence of missing mass: a yet-unseen type of matter distributed throughout the universe whose presence could explain several problematic astronomical observations. These papers indicate the establishment of the 'dark matter' problem, one of the most well-known anomalies in the currently prevailing cosmological model. This model states that 85% of the universe's mass budget consists of dark matter, but that its nature is yet to be discovered. To date, the problem of dark matter has not been solved. With high-energy particle collider experiments and multi-wavelength astronomical observations, researchers are probing an immense parameter space to find out what is the nature of the dark matter, but despite these efforts, the dark matter hunt has produced more than four decades of null-results.

In this paper, I use the historical establishment of the dark matter problem as a case to study how anomalies form in scientific practice. Although the establishment of the dark matter problem traces back to 1974, the two observations on which the papers base their conclusion had been around for much longer. In the 1960s the radial constancy of galaxies' rotational velocity was observed, and as early as in the 1930s it was known that the masses of galaxies did not add up to make sense of the dynamical behaviour of clusters of galaxies. Both observations are in hindsight considered evidence for the existence of dark matter, but only in 1974 these results were put together as a single consistent problem. By addressing the conditions of dark matter's establishment at that time, I illustrate the process of how certain scientific results become problematic. Specifically, I trace how the readily available astronomical observations transformed into evidence for missing matter by illuminating two aspects of this history: the translation of observations into a new disciplinary context, and the argumentative means by which this was done.

I show how the two hybrid research groups, consisting of astronomers and physicists, translated existing research from different subfields in astronomy to a new type of astrophysical practice that arose in the 1960s: physical cosmology. Through 'a priori' cosmological assumptions on the mass density of the universe, and an inference to a common origin, the authors of the paper realised an epistemic transformation in which two observed phenomena turned into evidence for a single problem. I argue that in this transformation, an anomaly was formed by the retrospective recognition of observations as evidence. By studying how dark matter was made to matter I then hope to elucidate the contingent ways evidence is used, argumentation takes place, and anomalies form in astrophysical practice.

Symposium: Evaluation, Quality and Success in Interdisciplinary Research

Organizers: Jaana Eigi, Endla Lõhkivi.

Contributors:

Jaana Eigi, Katrin Velbaum, Endla Lõhkivi, Edit Talpsepp-Randla & Kristin Kokkov.

Inkeri Koskinen.

Hauke Riesch.

From Thomas Kuhn's (1996; first edition 1962) account of scientific communities and paradigms guiding them to Helen Longino's (1990) argument about the objectivity-maintaining role of transformative criticism in communities that share norms and avenues for criticism, philosophers of science have been showing the importance of communities for the creation and evaluation of knowledge claims.

With the rise of interdisciplinary research, however, a community where knowledge is produced may no longer be expected to be united in any such way, as representatives of various disciplines as well as non-scientists work together in inter- and transdisciplinary projects.

Given the importance of the shared background for epistemic practices in research communities, this development raises a number of interesting questions. What kind of obstacles do researchers face in interdisciplinary research and how much of interdisciplinarity they are able to achieve? Do participants of interdisciplinary projects experience problems related to evaluating other participants' contributions, having their own contributions evaluated by other participants or their overall project being evaluated from the outside? How does interdisciplinary work influence one's self-conception

and the perception of one's competence as a researcher? Do interdisciplinary and transdisciplinary projects offer new robust criteria for evaluation thanks to, for example, their greater practical impact? Or does their practical successfulness raise new problems for evaluation? If interdisciplinary research may indeed be expected to face certain problems, why does it remain attractive and what does it say about our understanding of traditional disciplines?

The presentations in this symposium address these questions using a variety of approaches, including an analysis of reasoning in an interdisciplinary field and two qualitative studies building on the interviews with interdisciplinary researchers in an established interdisciplinary field and with researchers and extra-academic collaborators in a transdisciplinary project. Together, they demonstrate how paying philosophical attention to different aspects of interdisciplinary practices helps to understand the current reality and the potential possibilities of interdisciplinary research.

Kuhn, Thomas (1996). *The Structure of Scientific Revolutions*. Chicago; London: University of Chicago Press (3rd edition).

Longino, Helen (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press.

Misunderstandings and Epistemic Misjudgements in an Interdisciplinary Field and How Researchers Live with Them

Jaana Eigi – University of Tartu, Estonia

Katrin Velbaum – University of Tartu, Estonia

Endla Lõhkivi – University of Tartu, Estonia

Edit Talpsepp-Randla – University of Tartu, Estonia

Kristin Kokkov – University of Tartu, Estonia

There seems to be no doubt that successful, productive, well-functioning and well-regarded interdisciplinary research areas exist – computational linguistics and language technology are some examples. Yet thinking about the very possibility of successful interdisciplinary research may bring some genuine puzzlement. How can one make sense (and recognise the quality) of knowledge claims made in a different discipline when one lacks knowledge, skills and the understanding of aims that are a given for the members of that discipline? How can one make oneself understood and appear credible in the eyes of representatives of other disciplines under such conditions?

The aim of the presentation is to explore some ways these issues may emerge as well as ways they may be resolved. The presentation does so by analysing

six semi-structured interviews conducted with researchers in computational linguistics and language technology.

The presentation demonstrates how misunderstandings and misjudgements of quality may and do arise in these areas of interdisciplinary research. As a result, there may be a failure to appreciate work done by partners from a different discipline; or a failure to use its results in a specific research project; or a failure to see its importance for the development of the interdisciplinary field more generally. There may also be misjudgements of the quality or the importance of a discipline's contribution on a more global level, manifesting themselves, for example, as insufficient funding for specific directions in interdisciplinary research.

The presentation also demonstrates how researchers may resolve these issues or cope with them in some alternative way. It describes a variety of strategies for ensuring some shared ground with representatives of other disciplines or for working together in the absence of such a shared ground. It also shows how researchers use a variety of criteria to judge the quality of research they produce in the interdisciplinary field, from the mutual relevance with the results of research done in other, more traditional fields, to successfully working practical applications, and to such metrics as the number of publications, conferences and defended theses in the interdisciplinary field.

While the evaluation of one's own and others' research work is the central topic of the presentation, it also shows how these issues, and their resolution, are closely related to these researchers' self-understanding as interdisciplinary researchers and their vision what interdisciplinarity means in their field.

The presentation thus aims to use empirical interview material in order to explore actual interdisciplinary practice in a field. It shows how even in a well-established and successful field there may be a variety of problems related to the understanding of the quality of one's and others' research as well as a variety of ways to resolve them, to work around them or to live with them when doing interdisciplinary research.

Epistemic success and societal impact in extra-academic collaboration

Inkeri Koskinen – University of Helsinki, Finland

Collaboration with extra-academic agents is nowadays common in science. Especially when the aim is to produce practically usable knowledge, and solve pressing problems, stakeholders and extra-academic experts are included in research teams. Various forms of collaborations are being developed in diverse fields; they range from co-research with private enterprises to activist research

initiated by stakeholder groups. They however share one goal: increasing the societal impact of academic research.

Philosophers, historians, and sociologists of science have examined cases of successful collaborations across disciplinary boundaries and across the boundaries of science. But this literature usually presupposes that success in such collaborations depends on whether the collaboration succeeds epistemically: whether epistemic exchange takes place, whether new findings are made, methods developed, etc.

In science policy, however, success in extra-academic collaboration is often taken to mean success in creating societal impact: solutions to practical problems, commercializable products, policy-relevant results. There seems to be an implicit assumption that a collaboration that fails from an epistemic point of view, cannot succeed in creating beneficial societal impact. I question this assumption.

I illustrate my claim with a case study: I have followed a 2-year project in which the research team consisted of sociologists, artists, and journalists. I attended their research meetings and interviewed all participants. As in many phases of the project the subgroups worked quite independently of each other, I focused on two collaborative phases: in the first the sociologists collaborated with the journalists, and in the second, with the artists.

From an epistemic viewpoint, the collaboration between the sociologists and the journalists succeeded: by conducting a survey in a major newspaper they created a boundary object that produced data for the sociologists and was a source of several articles for the journalists. Together they were also able to create a solution to a methodological problem that troubled the sociologists.

Considered from the same viewpoint, the collaboration between the sociologists and the artists largely failed. Many of the initial objectives were abandoned, the sociologists and the artists never agreed even on the starting points of the collaboration, and finally the subgroups worked independently without much epistemic exchange taking place.

However, if the criterion of success is taken to be the created societal impact, both collaborations succeeded. Both created more public interest in the work of the whole group than would have been likely without the collaboration, and this interest led to policy outcomes. So the societal impact of an extra-academic collaboration does not necessarily depend on whether the collaboration succeeds epistemically or not.

I argue that to understand the relationship between epistemic success and success in creating societal impact in extra-academic collaborations, it is necessary to differentiate between different types of societal impact. I then conclude by discussing the possibility of situations in which a collaborative project produces epistemically dubious results but succeeds in creating the wanted societal impact. If such situations are indeed possible, it is particularly

important to recognise the looseness of the link between epistemic success and societal impact.

Inter-Discipline and Punish

Hauke Riesch – Brunel University London, United Kingdom

Interdisciplinarity is, like excellence (Moore et al. 2017), commonly acknowledged as a “good thing” in universities, but like many good things, a precise definition and a completely convincing argument for why it is a good thing has not yet been forthcoming (Jacobs and Frickel 2009). First, I will look at the potential dynamics of interdisciplinary groupings, identifying four types of interdisciplinarity based on the interactions of disciplines viewed as complex social identities (Riesch 2014). Second, I will argue that discourses on interdisciplinarity often presuppose rarely acknowledged assumptions about disciplines and their functions, and that interdisciplinary approaches need to take into consideration the various natures of the disciplines that are meant to combine: interdisciplinarity may have to be handled differently depending on the disciplines in question, and we may also need to consider what we might potentially lose through abandoning disciplines, or whether by combining them we don't simply produce new disciplinary spaces that suffer from the same (possibly imagined) shortcomings as the previous divisions within the academy (Riesch, Emmerich and Wainwright, under review). Combining these perspectives, the discourse surrounding interdisciplinarity will be analysed as what Prainsack and I (2016) have called a “fantasy of redemption”- a useful rhetorical space in contrast to which any real or imagined failures of science and academic research can be packaged as in need of salvation. For these discursive purposes the vagueness of the concept is its strength: an intellectual Rorschach test that can be whatever is required to save science from its problems.

Jacobs, J. A., and Frickel, S. (2009). Interdisciplinarity: A critical assessment. *Annual review of Sociology*, 35, 43-65.

Moore, S., Neylon, C., Eve, M. P., O'Donnell, D., and Pattinson, D. (2017). Excellence R Us: University Research and the Fetishisation of Excellence. *Palgrave Communications* 3, p. 16105

Prainsack, B. and Riesch, H. (2016). 'Interdisciplinarity Reloaded? Drawing lessons from “Citizen Science”', in Frickel, S., Albert, M. and Prainsack, B. (eds.) *Investigating Interdisciplinary Collaboration: Theory and Practice across Disciplines*. Rutgers University Press

Riesch, H. (2014). 'Philosophy, history and sociology of science: Interdisciplinary relations and complex social identities'. *Studies in History and Philosophy of Science Part A*, 48, 30-37

Riesch, H.; Emmerich, N. and Wainwright, S. (under review) "Introduction: Crossing the divides" in Riesch, H., Emmerich, N. and Wainwright, S. (eds): *Crossing the divides: Sociologies and Philosophies of Bioethics*. Springer

Symposium: Causal and Informational Specificity in Biological Practice: Unchallenged Assumptions and Neglected Dimensions

Organizers: María Ferreira Ruiz, Alan C. Love, Janella Baxter, Oliver Lean, Catherine Kendig.

Contributors:

Janelle Baxter.

María Ferreira Ruiz.

Oliver Lean.

Alan Love.

Philosophy of biology has become increasingly interested in the concept of specificity. This interest arose amid the controversy about genetic determinism—about whether, and to what extent, an organism's traits are determined by their genes. After the "interactionist consensus" (Kitcher 2001), some retain the intuition that genes have, nevertheless, a privileged ontological, investigative and/or explanatory role. The current state of the debate revolves around whether or not the concept of specificity can cash out this intuition.

Philosophical discussions of specificity come in two flavors: informational and causal. Informational specificity has roots in Crick's "central dogma of molecular biology," according to which "information" broadly referred to the precise determination or specification of a protein sequence by a corresponding sequence of DNA (mediated by a complementary RNA sequence). The complexity of eukaryote genetics has led some to argue for an informational parity between genes and other factors (Griffiths and Gray 1994; Sterelny et al. 1996; Oyama 1985; ; Griffiths and Stotz 2013). Causal specificity has its roots in Lewis' notion of influence (2000) and was developed by Woodward (2010) to account for the degree of "fine-grained control" a cause has over its effect. Causal specificity has been cited as an important criterion of causal selection, and some claim that while genes do not operate alone, they specify their developmental outcomes to a high degree (Waters 2007; Weber forthcoming, 2017; Stegmann 2012). Recently, it has also been invoked to

support the opposite thesis that genetic and non-genetic factors are causally on a par. Griffiths et al. (2015) show that the specificity of a causal relationship can be quantified as the mutual information between the variables representing cause and effect, and the incorporation of communication theory into the debate links the informational and causal senses of specificity.

Because of its origins, the discussion of specificity has become developmentally entrenched. This has either obscured or left unaddressed diverse conceptual and practical issues. Focus on the causal specificity that genes and non-genetic causes have on the amino acid sequences of proteins not only suggests that switch-like causes are less important to biological explanation, but it also overlooks the question of whether the philosopher's use of "specificity" properly reflects molecular biologists' use of the term. Additionally, recent attempts to analyze informational talk in biology in terms of causal specificity and the mathematical theory of communication assume that the main concerns regarding informational talk in biology are thereby overcome. Moreover, it is assumed that the amount of information a variable carries is significant per se, regardless of the investigative context.

This symposium turns to biological practice to grapple with several neglected dimensions and unchallenged assumptions surrounding causal and informational specificity: Is control actually all about the specificity of a causal relation? Does specificity imply the privileging of causes at all? Does a specificity approach really enable a substantive account of biological information? And, does causal specificity—as discussed by philosophers—really capture the notion of specificity as used in molecular biology?

Griffiths Paul, and Russell Gray. 1994. "Developmental systems and evolutionary explanation." *The Journal of Philosophy* 91:277–304.

Griffiths, Paul, Arnaud Pocheville, Brett Calcott, Karola Stotz, Hyunyu Kim, and Rob Knight. 2015. "Measuring causal specificity." *Philosophy of Science* 82:529–555.

Griffiths, Paul, and Karola Stotz. 2013. *Genetics and Philosophy: An Introduction*. New York: Cambridge University Press.

Kitcher, Philip. 2001. "Battling the undead: How (and how not) to resist genetic determinism." In Rama S. Singh, Costas B. Krimbas, Diane B. Paul, & John Beatty (Eds.), *Evolution: Historical, Philosophical and Political Perspectives* (Festschrift for Richard Lewontin). Cambridge: Cambridge University Press (pp. 396–414).

Lewis, David. 2000. "Causation as influence." *The Journal of Philosophy* 97:182–197.

Oyama, Susan. 1985. *The Ontogeny of Information: Developmental Systems and Evolution*. Cambridge: Cambridge University Press.

- Stegmann, Ulrich. 2012. "Causal control and genetic causation." *Noûs* 48:450–465.
- Sterelny, Kim, Kelly Smith, and Michael Dickison. 1996. "The extended replicator." *Biology & Philosophy* 11:377–403.
- Stotz, Karola, and Paul Griffiths. 2017. "Biological Information, causality and specificity – an intimate relationship." In Sara Imari Walker, Paul Davies, and George Ellis (eds.), *From Matter to Life: Information and Causality*. Cambridge: Cambridge University Press (pp. 366–390).
- Waters, C. Kenneth. 2007. "Causes that make a difference." *The Journal of Philosophy* 104:551–579.
- Weber, Marcel. forthcoming. "Causal selection vs. causal parity in biology: relevant counterfactuals and biologically normal interventions." In C. Kenneth Waters and James Woodward (eds), *Causal Reasoning in Biology*. Minneapolis: University of Minnesota Press.
- . 2017. "Discussion note: which kind of causal specificity matters biologically?" *Philosophy of Science* 84:574–585.
- Woodward, James. 2010. "Causation in biology: stability, specificity, and the choice of levels of explanation." *Biology & Philosophy* 25:287–318.

Causation and Information in Biology – Arguments for Eliminativism

María Ferreira Ruiz – University of Geneva, Switzerland

Informational language is pervasive in biological practice: genes are customarily regarded as informational molecules and genetic mechanisms are described in terms of information being transcribed, translated, edited, and copied. Yet, the use of informational language in biology has been questioned and challenged (Oyama 1985; Maynard Smith 2000; Griffiths 2001). In philosophy of biology, the problem is generally presented as one of literal versus metaphorical use of language—whether DNA literally carries information or this is just a manner of speaking. Philosophers favoring the former view dedicate efforts to explicating this purported literal meaning, offering refined conceptualizations of information under which DNA can be said to properly carry information. Several accounts of this type have been proposed, but they have received diverse criticisms and there is no consensus about whether a best approach exists for explicating the informational talk in biological discourse.

A recent approach argues that most informational talk in biology is nothing but specificity talk (Griffiths et al. 2015; Stotz and Griffiths 2017; Griffiths 2017; Calcott et al. forthcoming; Pocheville et al. forthcoming). On this view, specificity is analyzed as fine-grained influence and then shown to be

measurable by means of the standard formalism of information theory. Against claims that, in biology, ‘information’ is meaningless, only a metaphor, or cannot be rigorously accounted for, proponents of the specificity account of biological information (SAI) set out to provide a robust, substantive concept of information, and argue that it generalizes to entities other than genes, thereby providing a theory of biological information.

I contend that SAI—however interesting and fruitful as a formal framework for causal analysis—provides an argument for the elimination of informational talk in biology rather than a robust, substantive account of its nature. First, even if much of the informational talk in biology refers to nothing but causal relations of high specificity, this does not tell us why it is correct to apply the concept of information apart from mere customary use. It might be objected that the basis for SAI’s claim to solve the problem of biological information is its use of information theory, the theory plays no clarificatory or “substance-giving” role in the account. Rather, it is used in an instrumental manner that is also manifest in ecological practice, where information theory is used for measuring species diversity in a given community (Begon et al. 2006). Second, if an account of informational talk in biology succeeds in explicating this notion by showing that information is nothing but specific causation, then this proves—without further arguments—that we can do without such informational talk. Ultimately, the key questions about the invocation of information in biological investigation and explanation remain open.

In examining what a substantive notion of information for biology would be, I conclude that more attention is needed to clarifying the structure of the relevant philosophical problem and articulating the criteria of adequacy that must be met for any account of biological information to be satisfactory.

Begon, Michael, Colin Townsend, and John Harper. 2006. *Ecology. From individuals to ecosystems* (4th. edition). Oxford: Blackwell Publishing.

Calcott, Brett, Paul Griffiths, and Arnaud Pocheville. Forthcoming. “Signals that make a difference”. *British Journal Philosophy for the Philosophy of Science*: <https://doi.org/10.1093/bjps/axx022>.

Griffiths, Paul, Arnaud Pocheville, Brett Calcott, Karola Stotz, Hyunyu Kim, and Rob Knight. 2015. “Measuring causal specificity.” *Philosophy of Science* 82:529–555.

Griffiths, Paul. 2017. “Genetic, epigenetic and exogenetic information in development and evolution.” *Interface Focus* 5:1–8.

——— 2001. “Genetic information: a metaphor in search of a theory.” *Philosophy of Science* 68:394–412.

Maynard Smith, John. 2000. “The concept of information in biology.” *Philosophy of Science* 67:177–194.

Oyama, Susan. 1985. *The Ontogeny of Information: Developmental Systems and Evolution*. Cambridge: Cambridge University Press.

Pocheville, Arnaud, Maël Montévil, Karola Stotz, and Paul Griffiths. Forthcoming. "Crick information: giving substance to biological information" (unpublished manuscript).

Stotz, Karola, and Paul Griffiths. 2017. "Biological Information, causality and specificity – an intimate relationship." In Sara Imari Walker, Paul Davies, and George Ellis (eds.), *From Matter to Life: Information and Causality*. Cambridge: Cambridge University Press (pp. 366–390).

When INF-Specificity is too Much of a Good Thing: A Defense of Switch-like Causation

Janella Baxter – University of Minnesota, United States

Several authors (Waters 2007; Stegmann 2012; Weber 2013, 2017; Griffiths et al. 2015) have embraced Woodward's (2010) account of INF-specificity as an analysis of the molecular mechanisms that determine the amino acid sequences of proteins. These authors regard INF-specificity as a privileged type of explanatory property over switch-like causal control. Switch-like causal variables are like the "on/off" button of a radio; whereas, causes with INF-specificity are like the radio station dial. While a switch-like cause can only take one of two possible values, causes with INF-specificity can take a range of alternative "settings," each of which systematically associates with one and only one "setting" of an effect variable. Thus, INF causal variables provide numerous opportunities to manipulate and control the value of an effect variable. Moreover, they can provide numerous answers to counterfactual questions about how the value of an effect variable will change given changes in the value of the causal variable. Commenters involved in the debate about INF-specificity accord a significant explanatory status to INF causal variables because of the greater amount of control and explanatory power they possess.

Philosophical focus on INF-specificity, thus, may be taken to suggest that switch-like causes are generally less important in biological explanations. And yet, switch-like causal structure characterizes much of the biological world. As a consequence, switch-like causes feature prominently in particular explanatory contexts. Switch-like causes are often crucial for regulating the initiation of irreversible developmental and cellular stages by translating continuous processes into an all-or-nothing affair (Yao et al. 2008; Capel 2017). Furthermore, switch-like causal control is often a powerful experimental approach for probing the actual activities and biological roles that molecular

components (including genes) perform across biological systems (Housden et al. 2017).

This talk articulates the significance of switch-like causal variables to account for why biologists sometimes privilege them over causes with INF-specificity in particular contexts of explanation and experimentation. By using Woodward's (2010) proportionality constraints for determining the appropriate "fit" between cause and effect variables, I will argue that in some cases it makes little conceptual and empirical sense to describe some types of biological causes in more fine-grained ways. Furthermore, manipulation of switch-like causes sometimes have experimental virtues that manipulation of INF causal variables don't. The amount of causal control seemingly available from INF-specificity frequently leads to inefficient experimental results due to noise and risks creating novel effect outcomes that would otherwise be absent in the biological system of interest.

Capel, Blanche. 2017. "Vertebrate sex determination: evolutionary plasticity of a fundamental switch." *Nature Reviews Genetics* 18:675–689.

Griffiths, Paul, Arnaud Pocheville, Brett Calcott, Karola Stotz, Hyunyu Kim, and Rob Knight. 2015. "Measuring causal specificity." *Philosophy of Science* 82:529–555.

Housden, Benjamin E., Matthias Muhar, Matthew Gemberling, Charles A. Gersbach, Didier Y.R. Stainier, Geraldine Seydouxx, Stephanie E. Mohr, Johannes Zuber, and Norbert Perrimon. 2017. "Loss-of-function genetics tools for animal models: cross-species and cross-platform differences." *Nature Reviews* 18:24–40.

Stegmann, Ulrich. 2012. "Causal control and genetic causation." *Noûs* 48:450–465.

Waters, C. Kenneth. 2007. "Causes that make a difference." *The Journal of Philosophy* 104:551–579.

Weber, Marcel. forthcoming. "Causal selection vs. causal parity in biology: relevant counterfactuals and biologically normal interventions." In C. Kenneth Waters and James Woodward (eds), *Causal Reasoning in Biology*. Minneapolis: University of Minnesota Press.

———. 2017. "Discussion note: which kind of causal specificity matters biologically?" *Philosophy of Science* 84:574–585.

Woodward, James. 2010. "Causation in biology: stability, specificity, and the choice of levels of explanation." *Biology & Philosophy* 25:287–318

Yao, Guang, Tae Jun Lee, Seiichi Mori, Joseph Nevins, and Lingchong You. 2008. "A bistable Rb-E2F switch underlies the restriction point." *Nature Cell Biology* 10:476–482.

Chemical Specificity is not Fine-Grained Control

Oliver Lean – University of Calgary, Canada

Chemical specificity is what allows biological molecules like proteins and DNA to interact with only one or a few chemical species. Specificity comes from a combination of two kinds of relation between interacting entities: physical fit between the surfaces (stereochemistry), and complementary electrical charges (electrochemistry). The result is that the activity of biomolecules can be highly targeted. Chemical specificity is critical for the precise organization of biological processes in space and time.

Specificity is also currently at the centre of a prominent debate in philosophy of biology. This debate is a successor to the controversy about genetic causation — whether genes play some unique or special role in development relative to other factors. This debate is currently about whether genes are relatively causally specific in the sense that they offer “fine-grained control” over their products, while other factors are like on/off switches with only coarse-grained effects (e.g. Waters 2007). This fine-grained specificity has gained close attention especially since (Woodward 2010), and has more recently been given rigorous quantitative treatment (Griffiths et al. 2015).

One might be forgiven for thinking that the notion of “specificity” at work here is the same kind that is central to molecular biology. However, chemical specificity is critical even in mechanisms that are switch-like, such as hormones activating cell receptors, and these are the definitional opposite of fine-grained specificity. Instead, perhaps chemical specificity is better captured by the notion of single causes producing single effects. Woodward (2010) discusses both this “one-one” specificity and the fine-grained kind, and suggests some connections between them. However, since chemical specificity is implicated in processes exemplifying both of these causal notions, it cannot be straightforwardly analyzed as either one or the other. The relationship between the molecular biological and causal notions of specificity is more complex, and in need of philosophical analysis.

This paper discusses chemical specificity and its relation to causal notions of specificity. Both fine-grained and one-one causal specificity are criteria of causal selection — features that make those causes interesting, relevant, or desirable in a given practical context. But they are also relatively abstract features: two causal relations in very different scientific domains — e.g. molecular genetics and social science — can share them. My central claim is that chemical specificity, as it is used in molecular biological practice, is not an instance of either kind of causal specificity; however, it is an explanation of how causal specificity is achieved in the concrete case of biomolecules. This fits

with the term's history: from its inception, molecular biology aimed to discover the mechanisms by which "biological specificity" was inherited. The elucidation of DNA structure and its chemically specific base pairing was considered the answer to this problem. It also fits with the concept's current use: in drug design, for instance, the aim is to develop drugs that are causally specific (e.g. with few side effects), while chemical specificity is seen as a means by which this end can be achieved.

Griffiths, Paul et al. (2015) "Measuring causal specificity." *Philosophy of Science* 82:529–555

Waters, C. Kenneth (2007) "Causes that make a difference." *The Journal of Philosophy* 104:551–579

Woodward, James (2010) "Causation in biology: stability, specificity, and the choice of levels of explanation." *Biology & Philosophy* 25:287–318

Positional Information and the Measurement of Specificity

Alan Love – University of Minnesota, United States

Philosophers have long compared the relative importance of genes to other causes and many analyses have concentrated on the concept of specificity as fine-grained influence in relation to difference-making accounts of causal explanation. Recently, a quantitative measure of causal specificity using information theory has been offered as a means to arrive at more accurate comparisons of different causes (Griffiths et al. 2015). However, competing accounts of the amount of information available in a gene sequence compared with the process of alternative splicing remain (Weber 2017).

These appeals to information and specificity are intended to substantiate or undermine the privileging of genetic causation in biological investigation and explanation. Yet little effort has been expended on understanding practices where scientists measure specificity in biological systems. Here I scrutinize an example of this type of practice: measuring positional information in embryonic patterning of *Drosophila*.

Positional information is the idea that cells acquire an identity in relation to their relative position in a bounded region of the embryo. Once acquired, cells differentiate into distinct types based on characteristic patterns of gene expression, which facilitates the formation of higher-level patterns, such as repeating stripes of cells that eventually yield morphological segments. Dubuis and colleagues (2013) measured positional information in the expression of several genes at one stage during *Drosophila* embryogenesis. They found that the corresponding amount of information is enough to specify the location of

cells along the anterior-posterior axis of the embryo within an error rate of 1% and corresponds closely to what is needed for each cell to have a unique identity.

These biologists did not compare the amount of information carried by genes with that of other factors or calculate the information carried by a particular gene. Instead, they quantified the amount of information in a gene expression pattern throughout the embryo at a time to understand how cells differentiate and produce morphological patterns. The relevant measure of mutual information is not of a DNA sequence relative to a protein sequence, but of an amount of gene expression relative to a location along the axis of the embryo. This example speaks to various proposals regarding causal specificity. For example, the actual values that a causal variable takes in a population were favored over the range of all possible values or some restricted range of relevant values. More generally, these reasoning practices demonstrate that the measurement of information is relativized to a biological question at a stage in the life history of the organism. Biologists are less concerned with how much information is contained in a factor and focused on whether the measured specificity explains the phenomenon under scrutiny.

- Dubuis, J.O., G. Tkacik, E.F. Wieschaus, T. Gregor, and W. Bialek. 2013. Positional information, in bits. *Proceedings of the National Academy of Sciences USA* 110:16301–16308.
- Griffiths, P., A. Pocheville, B. Calcott, K. Stotz, H. Kim, and R. Knight. 2015. Measuring causal specificity. *Philosophy of Science* 82:529–555.
- Weber, M. 2017. Which kind of causal specificity matters biologically? *Philosophy of Science* 84:574–585.

Symposium: The Changing Nature of Mathematical Solutions How Computer Methods Affect the Concept of Mathematical Solutions

Organizers: Nicolas Fillion, Johannes Lenhard.

Contributors:

Mattias Brandl & Johannes Lenhard.

Nicolas Fillion & Jabel Ramirez.

Julie Jebeile & Vincent Ardourel.

Robert H.C. Moir.

The concept of solution to mathematical problems is central in science, because many scientific problems are solved with the help of mathematics. This concept, however, is even more important, because it serves as a

paradigm (or ideal case) for what solving a problem means. The nature of a mathematical solution is of utmost clarity and distinction: A term solves an equation iff this equation holds when the term is inserted. Moreover, such solutions are of great epistemic value. The mathematical equations in many theories and models implicitly describe how certain input and output variables are related, a relation that a solution makes explicit. For instance, a closed-form solution describes how the resulting output varies as a well-understood function of the input. Knowing the solution hence enables predictions of phenomena, explanation of the behavior of the modelled system, and other crucial kinds of scientific inferences.

Modern applied mathematics has importantly expanded this classical understanding of mathematical solutions. Using the computer as an instrument has widened the realm of mathematical modelling and theorizing greatly, mainly because finding a solution is possible in far more situations. A numerical solution relies on computational procedures that somehow approximate an analytical (ordinary, traditional) and solution, as do other related computational approximation methods. This symposium claims that this viewpoint is misleading. Numerical solutions do not simply extend the range where finding solutions is possible, but they affect and indeed transform the very concept of what a solution is. Since the concept of solution is a central one in science, changes and transformations of it are philosophically momentous.

The following aspects are examined by the contributions:

- “Numerical solution” is an umbrella term that in fact regroups a number of heterogeneous methods. This is because numerical solutions present compromises between conflicting criteria like tractability, speed, usability, and accuracy. How do scientists achieve a balance between them, and how does it give rise to different perspectives on mathematical solutions?
- Finding a numerical solution is a necessary condition for a sufficient understanding of mathematical models, at least in applied contexts. The directions into which models are developed are very different from the pre-computer time. Mathematical modeling is channeled in new ways that philosophy of science needs to discuss and investigate.
- The social and institutional organization of science changes, since finding a (numerical!) solution can often be outsourced to a software package or even to a computing center. These resources, however, come with their own conditions concerning how scientific problems have to be structured.

The Model Character of Solutions. A Challenge to Method and Authority

Matthias Brandl – Berlin, Germany

Johannes Lenhard – Bielefeld University, Germany

The concept of mathematical solution is central to quantitative sciences and it is momentous to an even wider domain of (non-quantitative) practices. Here is a definition in most simple and terms: “If a problem is formulated as an equation $f(x) = 0$, this problem is solved by a term a iff $f(a) = 0$.” We call this the “strict” concept of solution. However, the form of the function f might be complicated and hence finding a solution might require special competences in analytical, algebraic, or numerical methods. Having this competence proves professional authority on the side of the mathematical expert. This expertise includes the ability to describe (algorithmic) recipes for attaining a solution. The formula for quadratic equations (you all loved to use in school) is a pertinent example; those algorithms that are built into software packages are another.

This concept of mathematical solution has model character in an important sense. In many practical fields this concept serves a paradigm for what solving a problem means. Of course, in practice, problems usually require further conditions that complicate the case. Even if problems can be formulated mathematically, one often deals with approximations, not strict solutions. Or the problem cannot even be formulated in a fully quantitative fashion. Nevertheless, the strict concept remains significant as an ideal type that serves as a guide.

Most importantly, the concept of mathematical solution has model character also in fields that do not work with problems formulated in mathematical language. Professional guidelines for psychotherapists are a nice example to illustrate our viewpoint. They assume there is a problem that needs professional analysis and then can be solved based on this analysis. Method of analysis and solution are what determines the status and authority of the expert.

Our thesis has two parts. First, we argue that in mathematical practice a “weak” concept of solution is much more relevant than the strict concept. The concept of numerical solution is an instance of this weak sense. We will discuss so-called meta-heuristics, a cluster of numerical optimization methods. The solution (finding an optimum) there is characterized by being minimal (largely independent of problem), non-unique (many different solutions), and dependent on independent evaluation (not confirmed by analysis itself). All these properties stand in stark contrast to the strict concept of solution.

Second, we investigate into the reasons why the strict concept is serving as an ideal type in domains outside of mathematics. We find this role is unfounded since it rests on an overly rationalistic picture of mathematical practice. If fields like psychotherapy would accept that solutions are minimalistic, non-unique, and in need for independent evaluation, this would challenge not only the methodology but the interventions and authorities that are established in these fields.

A philosophical take on variational crimes

Nicolas Fillion – Simon Fraser University, Canada

Jabel Ramirez – University of La Laguna, Spain

In this paper, we discuss some epistemological innovations associated with the solution of mathematical problems by means of the finite element method. This method, used to obtain approximate solutions to partial differential equations within finite domains with possibly irregular boundary conditions, has received comparatively little attention in the philosophical literature, despite its efficiency handling complex real-world systems. This method with approximate solutions, contrary to exact solutions, involve error-control strategies. This is why assessing the validity of inexact solutions requires that we emphasize aspects of the relationship between solutions and mathematical structures that are not required to assess putative exact solutions. One such structural element is the sensitivity or robustness of solutions under perturbations, whose characterization leads to a deeper understanding of the behavior of the system. The transition from an epistemological understanding of exact solutions to the concept of an approximate solution can thus be characterized as structure enrichment. This transition generates an epistemological scheme to assess the justification of solutions that contains more semantic elements whose inner logic is essential to a philosophical understanding of the topic.

Examining other aspects of structure enrichment, we see that the mode of discretization of the domain of the finite element method and its interpolation of local solutions from the element nodes violates a number of principles considered mathematically and epistemologically essential for computational methods. And yet, the method has proved to be tremendously successful. Our objective is to characterize and contextualize the innovation and justification of the method in light of those differences with an ordinary view on mathematical solutions.

To be sure, there is a practical acceptance of the method by practitioners in order to overcome the representational and inferential opacity of the usual

numerical models. What makes it so advantageous to use this method in practice is its discretization scheme, which is applicable to objects of any shape and dimension. This innovative mode of discretization provides a simplified representation of the physical model by decomposing its elements into triangles, tetrahedral, mainly. Next, each element is locally associated with a piecewise low-degree polynomial that is interpolated to ensure sufficient continuity between the elements. On that basis, a recursive composition of all the elements is made to obtain the total solution. However, this presents a dilemma, since using piecewise polynomials that will be continuous enough to allow for a mathematically sound local-global “gluing” is often computationally intractable. Perhaps surprisingly, computational expediency is typically chosen over mathematical soundness. Strang has characterized this methodological gambit as a “variational crime.” We explain how committing variational crimes is a paradigmatic violation of epistemological principles that are typically used to make sense of approximation in applied mathematics. On that basis, we argue that the epistemological meaning of these innovations and difficulties in the justification of the relationship between the system and the solution lies in additional structural enrichment of the concept of validity of a solution.

Verification (& Validation) of Simulations against Holism

Julie Jebeile – Université catholique de Louvain, Belgium

Vincent Ardourel – KU Leuven, Belgium

The Duhem-Quine thesis states that a single theoretical hypothesis cannot be tested empirically in isolation, but all together with auxiliary hypotheses. The model-oriented version of this thesis has recently been addressed (e.g., Lenhard and Winsberg 2010; Winsberg 2010; Jebeile and Barberousse 2016; Lenhard 2018). What we shall call “Duhem problem” of refutation and confirmation holism states the following: when a model fails to match available data, it means that something must be wrong within the modeling assumptions, but the assumption(s) to blame cannot straightforwardly be identified.

In this paper, we focus on the specificity of Duhem problem when applied in the domain of applied mathematics. Here, in model validation, modeling assumptions are tested all together including those having a representational content (i.e., theoretical principles and simplifying hypotheses) and those related merely to the numerical scheme (i.e., discretisation of equations, meshing of the physical domain, and round-off). Therefore, “when a computational model fails to account for real data, we do not know whether to blame the underlying model or to blame the modeling assumptions used to

transform the underlying model into a computationally tractable algorithm” (Winsberg 2010, p. 24). A specific form of holism thus appears insofar as the mere computational aspects of the model implementation may interfere in the validation process with the representational content of the model, so that their respective contributions in the model performance cannot be assessed distinguishably.

We consider Verification & Validation (V&V), a methodology initially designed to legitimate simulation (Oberkampf, Trucano, Hirsch 2002) as a solution to this specific form of holism. Generally, verification is characterized as a mathematical problem, and validation as a physics problem. Verification aims to determine that the model is well implemented into algorithms and the equations are correctly solved, while validation aims to determine that the equations constitute an accurate representation of the target system for the purpose at hand. Thus verification is supposed to ensure that numerical errors related to the numerical scheme do not affect significantly the model outputs in the first place, before the model outputs be compared with empirical data in validation.

We first insist on that, for V&V to be a solution to holism, verification and validation must be separated and performed one after the other in this order. We then present arguments that rather support an entanglement between verification and validation (Winsberg 2010 and Lenhard 2018) as well as a mitigated claim (Morrison 2015). We finally argue that there is no specific form of holism in principle. There is nevertheless a specific form of holism in practice that can be overcome by degree depending on the requirement of the scientists. Our argument is gradual and based on the very many ways of processing verification—such as a priori justifications, method of manufactured solutions, and formal methods of verification—that have received less philosophical attention than validation so far.

Jebeile, J. & Barberousse, A. (2016) Empirical agreement in model validation, *Studies in History and Philosophy of Science Part A*, volume 56, 168–174.

Lenhard, J., & Winsberg, E. (2010). Holism, Entrenchment, and the Future of Climate Model Pluralism. *Studies in History and Philosophy of Science Part B*, 41(3):253–262.

Lenhard, J. (2018) Holism, or the Erosion of Modularity – a Methodological Challenge for Validation, to appear in *Philosophy of Science (PSA 2016)*.

Morrison, M. (2015) *Reconstructing Reality: Models, Mathematics, and Simulations*. Oxford Studies in the Philosophy of Science, USA: New York.

Oberkampf, W. L., Trucano, T. G., Hirsch, C. (2002) Verification, validation and predictive capacity in computational engineering and physics. *Applied Mechanics Review*, Volume 57, Issue 5, 345.

Winsberg, E. (2010) *Science in the Age of Computer Simulation*. University of Chicago Press.

Effective Logic: Stable Reasoning in the Presence of Error

Robert H.C. Moir – The University of Western Ontario, Canada

Reliable and efficient reasoning in science often requires the ability to find solutions to mathematical model problems, particularly as computational modelling becomes increasingly ubiquitous in scientific methodology. More often than not, however, it is not possible in practice to solve such model problems as originally posed, forcing transformations of the problem to make solution possible, transformations that typically introduce error.

Because error is introduced throughout the mathematical modelling process, since mathematical models succeed by focusing only on relevant details, there is good reason to regard so-called approximate solutions to mathematical modelling problems as bona fide solutions. This perspective can be made rigorous by considering input data, solutions, and problems themselves, in a context of variation. Then, any method that produces a solution to a (sufficiently) nearby problem can be understood as a (generalised) solution to the original problem. This approach can be rigorously justified in terms of concepts from a branch of modern error theory called backward error analysis. The purpose of this paper is to show how the production of generalised solutions to mathematical problems can be viewed as inferences in a generalised kind of logic, called effective logic, where inferences are stable under variation. The fundamental concept of this logic is effective validity, which is defined informally in terms of inferences where nearly true premises yield nearly true conclusions.

A natural consequence of the concept of effective validity is that inferences are only locally stable. This accords with the fact that modelling and computational methods that introduce error in order to produce solutions always have boundaries of validity, outside of which they fail to give valid solutions. Since such reasoning is the norm in scientific practice and everyday life, it is valuable to have an understanding of the basic logical structure of inference processes that are (locally) stable in the presence of error.

I will show how solution methods in science can be viewed as transformations of mathematical problems that facilitate making effectively valid inferences. The stability of the overall reasoning process can then be understood precisely in terms of a technical concept called inferential stability, which must obtain in any context where a mathematical model or computational solution method is used to describe, predict, explain or control a system or phenomenon.

Symposium: Data Practices in Personalized and Precision Medicine: How “Big Data Microscopes” Reconfigure Health and Disease

Organizer: Sara Green.

Contributors:

Stefano Canali.

Sophia Efstathiou.

Sara Green, Annamaria Carusi & Klaus Hoeyer.

Niccolò Tempini.

The possibility of combining omics data, health records and epidemiological data is currently promoted as a way to improve health care and biomedical research. The envisioned data-intensive efforts are referred to as personalized medicine, precision medicine, systems medicine, or even P4 medicine, as an acronym for medicine aimed to predict, prevent, personalize and make patients participate. In this symposium, we discuss prospects and challenges associated with a common goal of these streams: to improve the ability to understand and predict diseases through new ways of producing, sampling, and analyzing data.

Big data methodologies have been envisaged as a technological revolution akin to the invention of the telescope in physics or the microscope in the life sciences. The analogies and associated metaphor of “big data microscopes” may be far-fetched, but this framing raises philosophically interesting questions about the transformative potential of big data. To what extent is it possible to see the world in new ways through datasets that are unprecedented in their scope and level of detail? How do new difference-makers come to light through data analysis, and which aspects may be neglected as a result of the new focus? The metaphor of data microscopes, however, has severe limitations when it comes to understanding what it takes to make sense of data. The emphasis on enhanced vision may conceal how complex procedures for collecting, sampling, organizing and analysing data shape the constitution of evidence from population data. This symposium throws light on the complexity of these processes by examining the epistemic, social and organizational conditions for the development and application of the proposed strategies.

The symposium combines perspectives from philosophy of science and STS-research to reflect upon the following questions: How are “personalized” categories and characteristics identified through data analysis, and to what extent are the strategies dependent upon and constitutive of specific theoretical assumptions? Which traces of evidence are entailed when

attempting to provide data-intensive disease taxonomies? What are the epistemic and social implications of data-driven reconfigurations of disease and health? These questions will be explored through analyses of concrete examples. The first two talks examine how specific traits are legitimized as medically relevant differences through reconfigured evidence practices. Green, Carusi, and Hoeyer analyze epistemic and organizational challenges associated with a finer-grained disease taxonomy. Efstathiou examines how this challenge plays out in attempts to personalize RCTs, as exemplified in the African-American Heart Failure Trial. Through such practices, disease categories are reframed and redefined. In the third talk, Canali shows how other central notions in epidemiology, such as exposure and environment, undergo revision and renegotiation as a result of molecular technologies within the framework of exposome research. The fourth talk by Tempini explores the complex processes associated with data interpretation, integration, and reassignment of evidence in the context of the cancer genomics platform COSMIC and the self-reporting platform PatientsLikeMe. Together, the four talks provide contextualized accounts of the negotiation of evidence involved in the ‘personalization’ of medicine and the socially engrained data sources on which the techniques rely.

Molecular Data and Shifting Notions in Epidemiology

Stefano Canali – Leibniz Universität Hannover, Germany

Arguably, integrating and combining large quantities and different kinds of data is among the defining features of epidemiology. The addition of new sources of data is therefore considered a significant opportunity for the discipline. Molecular analyses are among the new sources of data that some projects in epidemiology are currently implementing. In this paper, I focus on projects applying the ‘exposome approach’ and consider the consequences that the use of molecular data has on the conceptual framework of epidemiology.

The main tenet of the exposome approach is that epidemiologists, when studying environmental risks associated with disease, should have a global attitude towards exposure, capable of distinguishing between internal and external components. In other words, exposure should not be reduced to environmental toxicants only: after these external elements get in contact with individuals, they may be present and initiate reactions within the organism, which should count as an internal component of exposure. In place of the traditional focus on external elements, the notion of the exposome is introduced as the total sum of exposures experienced by individuals, including

internal and external components. Omics data enters the picture as the way to study this internal layer of exposure: molecular analyses are used to look for biomarkers capable of tracing the presence of toxicants within the body. I claim that the use of molecular data as part of the exposome approach brings about changes and, possibly, tensions concerning the conceptual framework of epidemiology. Firstly, the exposome is far from being a settled notion, and two approaches have been put forward on how to interpret it. According to what is known as the “top-down approach”, the exposome should be limited to what can be measured at the molecular level, which can then be linked to exposure to environmental toxicants. On the other hand, the “bottom-up approach” maintains that the exposome should be measured by collecting exposure data on any aspect of the external environment. In the paper, I clarify the differences between these approaches and argue that the bottom-up approach can be considered the most novel theoretically, but may run the risk of producing “fishing expeditions”. Related to the interpretation of the exposome is a second tension, concerning the conception of environment and what is considered external and internal. The differentiation between internal and external components and the idea that to be exposed is to be exposed to both something external and internal have led to differentiations between an “internal” and an “external” environment and to consider the body as “an environment”. I argue that this shows how the boundaries of what is conceived as environment change in the exposome approach and influence how omics data is considered and used as health data.

Personalising RCTs: What is the right target?

Sophia Efstathiou – Norwegian University of Science and Technology, Norway

This paper investigates one central question: what randomized controlled trial (RCT) designs would legitimate population-specific drug efficacy claims? “Personalised medicine” or “precision medicine” are visions driving work in bioscience in Europe as elsewhere. Though big data is hoped to open up existing taxonomies of disease and to refine our understanding of human difference, this vision is building on existing classifications. On its way to becoming “personalised”, biomedicine is targeting population subgroups, including groups identified through race/ethnicity categories. To examine some of the challenges of population-specific pharmacogenetics or genomics, I examine the case of a drug developed to target a particular ‘race/ethnicity’ group in the US. This is the case of the African American Heart Failure Trial or A-HeFT [I-III].

A-HeFT was a randomized, double-blinded placebo-controlled clinical trial that tested a heart disease drug called BiDil on 1,050 people self-identified as African American. A-HeFT was terminated early because the treatment was so efficacious it was deemed unethical to keep withholding it from people on the placebo arm. Passing the trial led the FDA to grant its approval to BiDil (in June 2005) for its target, which made it the first drug to come out with a race-specific label on. So what was controversial about BiDil? A-HeFT demonstrated its efficacy on its target and emphatically so. What seems to have troubled researchers here was the selection of this target population as a target population to begin with. There was a great controversy in the science studies researchers studying the case (e.g. Sankar and Kahn 2005, Kahn 2013). The epistemological critique launched against BiDil can be (very roughly) summed up as follows. BiDil focused on efficacy among African-Americans, compared to placebo-treatment, but it did not demonstrate inefficiency for non-African-Americans. Without documented ineffectiveness among the complement of its target, it was perceived as controversial to target the trial and treatment to a specific population group.

Whether or not this critique is correct, the case brings up an interesting problem. Critical philosophical literature on RCT methodologies has examined the internal validity of RCT designs (e.g. the importance of randomisation, of blinding, of the placebo effect), or how to rank RCTs compared to other evidence-generating methods (e.g. historical controls and observation studies, or deduction from theory). However there has been little discussion of the possibility that selecting RCT target populations may compromise a trial's validity claims and its ethical and social value. What warrants the selection of a human subgroup as a clinical target? In the case of socially and historically identifiable race/ethnicity subgroups, biological and social scientific contributions to health compete for explanatory relevance but a social reality of other human subgroups may lay hidden in other cases.

Kahn, J. (2013). *Race in a bottle: The story of BiDil and racialized medicine in a post-genomic age*. Columbia University Press.

Sankar, P., & Kahn, J. D. (2005). BiDil: race medicine or race marketing? *Health Affairs*, October, 11.

Personalizing medicine? Reconfiguring conceptions of disease and evidence with population data

Sara Green – University of Copenhagen, Denmark

Annamaria Carusi – University of Sheffield, United Kingdom

Klaus Hoeyer – University of Copenhagen, Denmark

Personalized medicine promises to improve medicine through strategies that better account for variation between patients that influence disease development and treatment response. Through this focus, personalized medicine must confront an old and persistent problem in medicine, namely, how to account for the uniqueness of individual patients while meeting demands for statistical evidence. We examine this challenge in the context of calls for a more dynamic disease taxonomy, based on integration of health data and (gen)omics data. Disease taxonomies are continuously revised in response to developments in biomedical research and clinical practice. Recently, however, politically authorized reports have stressed the need for a new taxonomy that can accommodate insights from data-intensive research more rapidly. An explicit aim is to develop a taxonomy that allows for a finer-grained stratification of diseases and patient groups so as to provide the basis for more individualized treatments. We analyse what a revised taxonomy involves, both epistemologically and organizationally. We examine the epistemic assumptions underlying the reframing of difference as well as organizational requirements for new data-infrastructures. We argue that the pursuit of a more precise taxonomy not only reframes boundaries between disease categories, but also reconfigures what constitutes evidence. Personalized medicine promotes the epistemic ideal of diagnostic accuracy through a more precise stratification of treatment-response groups, biomarkers and risk factors. The aim is not only to increase the resolution of existing categories but also to revise these on the basis of available data and knowledge. In this process, the focus on specific differences or traits that make up the reference point in taxonomies may shift. This raises important questions about the relation between data and disease category, and on what basis the utility of new diagnostic tools can be evaluated. Without neglecting the potential for improved diagnostic categories and more targeted treatments, we stress that increased precision (in terms of increased resolution) need not result in increased diagnostic accuracy. The latter requires that the medical relevance of diagnostic markers, as well as the specificity and sensitivity of diagnostic tools, are established via reliable sources of evidence. We argue that meeting this demand for evidence is particularly difficult in the current context where phenotypic classifications, disease categories, the

validity of biomarkers, as well the causal understanding of disease are all under revision. Throughout the paper, we highlight how the epistemic challenges are aggravated in the current political context emphasizing accelerated speed of implementation. Given that it takes time to provide the desired evidence, we argue that it is appropriate to react with greater humility towards the challenges faced.

Making sense of precision data: the role of pathfinder processes and new information products in precision medicine infrastructures

Niccolò Tempini – Exeter, United Kingdom

Precision medicine aims at leveraging the flexibility that digital technologies allow (in storing, organising, configuring and processing increasingly complex assemblages of data and data structures). Its distinctive feature is to render descriptions of the world (entities, processes, procedures, therapeutic strategies) at a hitherto unforeseen level of granularity. For instance, it is increasingly affordable to sequence the entire genome of a healthy or a tumour cell. Genome sequencing can be used to investigate if specific features of its makeup are related to observed abnormalities or to suggest specific therapeutic strategies. In other settings, digital systems can allow patients to log comprehensive accounts of their symptoms and experience at very specific levels of description. With so much higher granularity it should become possible to discover razor-sharp causal relations and correlations, to be exploited or targeted through new therapeutic solutions.

In such a vast information space, researchers are confronted with a set of problems that can be defined as findability: what solutions would help and make available data interpretations, syntheses and references that are otherwise difficult to discover and use. This paper will demonstrate that the success of data infrastructures that are developed to support precision medicine depends on authoring new information products, constructing new data, repeatedly re-assessing available evidence. These processes are crucially aimed at changing the information ecology in which scientific practice is embedded, with a view to make newly available discoveries findable by other scientists. Actionability, the crux of precision medicine, is in this paper understood as crucially linked to successfully translating descriptions of the world, and the data associated to them. This translation moves away from the resolution at which precision data are generated, and towards levels of complexity that are better understood with the kinds of problem space that different users are working in (e.g. drug development angle vs diagnostic).

The paper builds on two case studies of precision medicine infrastructures: the cancer genomics platform COSMIC, and the participatory self-reporting platform called PatientsLikeMe. The paper identifies granularity and findability issues that are key to the successful adoption of the research infrastructures by the research communities they try to cater to, and the data reused by them. In the case of COSMIC, I show how this involves producing new data products through continuous re-evaluation and assessment of evidence, by redistributing, aggregating and decomposing several kinds of evidence claims. Similarly, in the context of PatientsLikeMe, a challenge is to make highly personal and self-reported accounts of patient experience and symptomatology usable as data for research. I show how this requires that a group of internal users continuously re-analyse the data to make different descriptions commensurable.

The paper observes how, in precision medicine infrastructure projects, internal data users are needed to take on a pathfinding role, thereby opening new avenues for the reuse of precision data. Processes of data reviewing, interpretation, editing and consolidation continuously evaluate the kinds of evidence claims that certain data should and should not support. They generate new data products such as new rankings, counts, classifications that are aimed at reflecting the resulting order of epistemic valuations, and consolidate it through their own mobilization.

Symposium: External Validity: What is it and how do we get it?

Organizers: María Jiménez-Buedo, Donal Khosrowi, Michiru Nagatsu, Federica Russo.

Contributors:

María Jiménez-Buedo.

Donal Khosrowi.

Michiru Nagatsu.

Federica Russo.

Over the last years, the distinction between internal and external validity has become an essential part of the conceptual toolkit for those engaged in the new experimental practices that have come to dominate many old and new subfields in both the biomedical and the social sciences. Thus, it is common to see references to the validity distinction when deciding whether to export the results of an RCT to a different population, when evaluating whether the results of an economic experiments are representative of the whole population behaviour, etc.

Although the notions, defined and redefined by Donald T. Campbell and a series of collaborators in the second half of last century, have a longstanding pedigree, their use was mostly restricted to social psychologists and related subfields until the terms were, at the turn of the century, picked up simultaneously by philosophers of science and by practitioners in emerging new experimental practices and the methodological discussion around them. In this way, the problem of external validity, has come to stand for a series of related crucial questions regarding extrapolation, generalizability of causal claims and even problems related to methodological artifacts linked to the ideal (but artificial) conditions found laboratory research.

The symposium tries to bring together several sets of issues related to the notion of external validity that have been recently subject to philosophical analysis. The aim is to find common grounds for reflection on a term that is often used ambiguously yet seems central to many philosophical discussions surrounding experimentation. The papers in the symposium will discuss both conceptual and methodological issues of external validity:

Regarding conceptual issues:

- The ambiguities implied in common uses of the terms internal and external validity and the exploration of different alternative meanings.
- The underlying assumptions embedded in the distinction between internal and external validity regarding the relationship between causal inference and concrete experimental designs.

Regarding methodological issues:

- Critical examination of the common assumption that there is a trade-off between internal and external validity in economics experiments. The relevance of this assumption is highlighted, and illustrated with examples from the literature on social preferences in both lab and fields behavioral experiments.
- Critical examination of the existing philosophical/methodological accounts of extrapolation. It is argued that the extrapolator's circle remains an unsolved problem for these accounts, and the desiderata for a better account are identified, with reference to examples from econometrics.

What is wrong (and right) with the distinction between internal and external validity?

María Jiménez-Buedo – UNED, Spain

The paper analyses the methodological implications of the ambiguous use of the notions of external and internal validity. The paper argues that at best, the terms serve the mere purpose of reminding us the importance of important

problems in experimentation, such as reliability or generalizability of findings. At worst the use internal and external validity can act as rhetorical devices and have the potential to supply ill-founded methodological advice. The paper explores the limited usefulness of the distinction between internal and external validity on the grounds of its numerous conceptual problems and recommends the use of viable, less problematic, terminological substitutes.

The paper and presentation are structured as follows:

Section 1 further explores the origins of the distinction between internal and external validity and analyses Campbell's purposes in coming up with these concepts. This introductory section traces how the distinction between internal and external validity gradually penetrated the philosophical literature on experiments and how the meanings of the terms gradually transformed from a list of potential threats associated to different well-known experimental designs (devised with an applied, pragmatic function in mind), to their more encompassing, philosophically loaded meanings today. The second section of the paper goes over some important conceptual problems of the internal/external validity distinction: first, we will go over the ambiguities in the common definitions of external validity. Second, we will go over some of the inconsistencies in the definitions of both external and internal validity and the object to which they are supposed to refer - be it experiments themselves (i.e., experimental designs), the data generated by the experiments, or the inferences from experiments. Favouring the latter interpretation, whereby internal and external validity would refer to the inferences from experiments, we present a critique, in the fourth section, of the assumptions implicit in the common uses of the distinction between internal and external validity. In particular, the paper argues that the distinction between internal and external validity presupposes an unrealistic conception of experimentation in which the relationship between inference and design is rigid, and underplays and misinterprets the role of background knowledge in the interpretation of experimental data.

Extrapolation of Causal Effects – Hopes, Assumptions and the Extrapolator's Circle

Donal Khosrowi – Durham University, United Kingdom

The extrapolation of causal effects from experimental populations to novel targets is a fundamental problem in Evidence-Based Policy (EBP) and econometrics. Populations often differ in causally relevant respects, so assuming that the effect of an intervention in the target is the same as in an experiment is typically not justified.

It's also clear that, at least in principle, information about similarities and differences between populations can be useful for extrapolation. Even if populations differ in important respects, learning how such differences bear on the effects of interest might still help us extrapolate successfully. At the same time, it will often be difficult to tell which similarities and differences matter most, and how information about them should be acquired and used.

Econometricians have recently developed interactive covariate-based approaches to address problems of extrapolation. (Muller, 2014, 2015, Hotz et al. 2005) They consider causally relevant differences between populations in the form of interactive covariates, i.e. variables that can induce differences in causal effects between individuals and populations. The aim of these approaches is to obtain a correct expectation of the effect of interest in the target by adjusting for differences in the distributions of such variables between populations.

I argue that a key problem with this idea is that learning what the relevant variables are for which to adjust, and how they are involved in producing the effects of interest will often be exceedingly demanding, e.g. one might need to know whether these variables play the same causal roles in both populations. That's not only burdensome but raises deeper concerns about the extrapolator's circle (Steel 2008): the causal knowledge about the target that's required may be so extensive that we could learn the effect we're interested in based on information about the target alone. This would render information from the experimental population redundant to our purposes.

As a potential remedy for this, I consider Daniel Steel's (2008) comparative process tracing (CPT) strategy that promises to help extrapolate while evading the extrapolator's circle. Although promising, I argue that CPT does not help evade the extrapolator's circle in many cases of interest in EBP and econometrics. To argue this, I offer a distinction between two kinds of extrapolation, attributive and predictive. Attributive extrapolation aims to causally attribute observed effects to their suspected causes. This means that both the intervention of interest as well as its suspected effects have been observed in the target. In contrast, predictive extrapolation aims to predict the future effects of yet unobserved interventions, so neither the intervention of interest nor its suspected effects have been observed in the target. I argue that this situation implies that the kind of evidence favored by econometricians, i.e. observational data from the target, cannot be used to tell whether populations are sufficiently similar to permit extrapolation.

I suggest that this provides strong reasons to think that EBP researchers and econometricians need to turn to alternative kinds of support for underwriting extrapolation, including mechanistic evidence, domain-general causal program theory and local knowledge of the target.

Analogical inferences as a solution to the extrapolator's circle: a case of field experiments movement in economics

Michiru Nagatsu – University of Helsinki, Finland

The main methodological challenge in extrapolating causal inferences that have been made using the experimental model system to the target system is called the extrapolator's circle: How can we establish the relevant similarity of the experimental model and the target system without already knowing enough about the latter? (Steel 2008) The main difficulty arises from the fact that the information concerning the target is scarce to support the relevant (i.e. useful for manipulation and prediction) analogy between the model and target. One prominent type of strategies, called comparative process tracing (Steel 2008, 88-92) is efficient use of non-experimental evidence, focusing on stages (i) of likely relevant differences between model and target given current background knowledge; and (ii) of the downstream mechanism, i.e., more direct causes of the outcome/effect. Less attention has been paid however to a more proactive type of strategies, which I call a generation of tailor-made evidence by experimentation. In short, this is an effort to consciously generate stronger evidence for analogical reasoning by designing experiments with the target in mind. Many examples can be found in the recent movement toward field experimentation in experimental economics. This movement is also often seen as a solution to the artificiality of the lab experiments due to excessive control, and the resulting limited extrapolability of causal inferences made in the lab to the field/wild/target. I will argue that it is unhelpful to see the problem of artificiality as a tension between internal and external validity, and to balance the two (this view is manifest in philosophical commentaries on experimental economics such as Woodward (2008)). This is because extrapolation by analogical inferences requires both types of validity. Methodologically, this implies that lab experiments, field experiments, and non-experimental field studies have to be conducted more systematically than currently practiced. The paper's contributions are two-fold: first it highlights a strategy to facilitate analogical reasoning that is alternative to much discussed comparative process tracing; second, in doing so with reference to the practice of field experimentation by economists, the paper provides its valid methodological rationale as well as practical recommendations to improve practice.

Extrapolation and the role of populational properties

Federica Russo – University of Amsterdam, Netherlands

Extrapolation, or external validity, is the inference by which results of one study, e.g. an experiment, are extended to a larger or a different population or to a different setting. Philosophers, in proposing an account of extrapolation based on the notion of mechanism, certainly do better than the dominant ‘statistically-minded’ approach of Cook and Campbell. However, mechanism-based extrapolators (i) do not recognise that there are external validity inferences other than the mechanism-based approach and (ii) fail to draw a distinction between ‘modelling’ and ‘using’ mechanisms. This distinction, I argue, is the step forward to undertake in order to strengthen the mechanism-based approach. In discussing the ‘modelling’ and ‘using’ mechanisms activities, I give reasons why mechanistic considerations ought to be used and I point to an aspect widely overlooked by all kind of extrapolators and that can, instead, help a great deal in both activities: the properties of the population under consideration. On the one hand, the reason why population properties, such as demographic or socio-political characteristics, play a key role in modelling the mechanisms is that they may determine different choices of variables or of proxies for some concepts, or they may determine the choice of the statistical model, and they may determine the interpretation of results. This is not to say, however, that explanation becomes a matter of subjective and arbitrary choices of the scientist, but it certainly makes explanation contextual. On the other hand, extrapolation involves much more than simply making considerations about the representativeness of the sample. Thus, the possibility to extrapolate strongly depends on the characteristics of the two populations undergoing comparison, and the process of extrapolation itself requires comparing the properties of the two populations—an aspect that was more apparent in external validity inferences of ethnographers than anywhere else.

Symposium: Narrative Practices in the Sciences

Organizer: Robert Meunier

Contributors:

Dominic Berry.

Robert Meunier.

Mary Morgan.

Matthew Paskins.

Scientists' accounts of their fields contain many narrative elements which can be found in various forms, both visual and linguistic, in scientific texts of all kinds. There are narratives about the subject matter of research as well as narratives about the research process, both of which can come neatly separated or be deeply intertwined; there are narratives about small facts and narratives concerning the big picture; finally, there are narratives that appear in published articles and narratives that figure in communication within smaller or larger, interdisciplinary groups of scientists. Such narratives reflect the material practices of the sciences in which they occur and constitute representational practices in themselves. For this reason, a focus on narratives can elucidate many topics pertinent to the philosophy of science in practice, from questions regarding the way experimental practice enters the discursive space, to issues concerning the integration of data from various research fields, to problems of case based reasoning. Is narrative a tool of inference, of explanation, of coherence-making, of prompting discovery? The session aims at addressing and exploring such questions about the role of narratives in the sciences, and participants will draw on various materials from the history and current developments in chemistry, embryology, biological engineering and anthropology.

The individual contributions thus build on and expand recent work on narratives. It has been argued that narratives are not limited to certain areas (e.g. history or historical sciences such as evolutionary biology), nor can their functions be reduced to ones of rhetoric nor the communication of scientific findings to broader audiences. Instead, they have epistemic and ontological functions associated with representation and explanation in many fields of the natural and social sciences (e.g. Morgan and Wise, 2017). While those studies have begun to develop an account of narratives in science, there are many so-far unasked questions about their roles, some of which are taken up in the session here proposed.

This session is in effect the first report of the new Narrative Science Project funded by the ERC and located at LSE. The papers here reflect the broad

scientific interests of team members ranging over the natural and social sciences, and their attempts to advance the philosophical analysis of narratives in the practices of science.

Supersaturated narratives: Data integration and coherence in biological engineering

Dominic Berry – London School of Economics and Political Science, United Kingdom

Why do some seeds fly, and what would count as a sufficient answer to this question? After all, one could generate an answer from evolution, or physiology, study of the environment, or indeed, an answer could be derived from engineering. My paper focuses on an ongoing collaboration between engineers and biologists dedicated to exploring and explaining how and why seeds fly, a project that sets itself the challenge of integrating insights from biology and engineering. The philosopher of science confronted with such a case might choose to mirror this programme and likewise integrate insights from the philosophy of biology and engineering, an approach pioneered by Bill Wimsatt. Building on these foundations, I focus on a broader level of abstraction, highlighting the central role of narrative for communication within the group, and subsequent successful data integration. The most appropriate way in which to understand this kind of collaboration is by attention to the shared narrative that the group is steadily building from various experiments and lines of inquiry. It is shown that one way in which scientists and engineers decide that data integration has been successfully achieved, is by arriving at a satisfying and 'supersaturated' narrative. Supersaturated because it has to incorporate the starting assumptions and interests of the whole group in its variety, regardless of their relevance for each individual.

This case matters for the philosophy of science in practice in three key ways. First, it brings attention to the intersection of biology and engineering, one that just like the intersection of biology and physics, or biology and chemistry, or biology and computer science, is important and worthy of exploration in its own right. Second, the paper highlights and defends the ubiquity and importance of narrative throughout the sciences and engineering. Where narrative is often reduced to 'story', or used as a pejorative label to suggest less reliable and less scientific knowledge, here narratives are defended on par with laws, inferences, deductions, and all those other candidate ways of knowing more commonly attended to by philosophers. Third, I add to the range of phenomena to which the philosopher of science in practice might attend when pursuing the question of data integration, one which has become

an important topic for the community. The paper is based on a series of interviews and short periods of laboratory ethnography that took place between 2015 and 2017. The work of scientists and engineers in practice is often constrained and directed by the ambition of producing a coherent and satisfying narrative. Coherence, an epistemological quality most recently emphasised by Hasok Chang, can be further explored by appreciation of its relations with narrative knowledge.

Ordering in research narratives and natural narratives

Robert Meunier – London School of Economics and Political Science, United Kingdom

The paper starts from the observation that in many sciences two types of narratives can be distinguished: 1) Narratives that concern the processes or entities and their properties and relations, as they are understood to happen or exist independently of their investigation. These might be called “natural narratives”. 2) “Research narratives” that concern the (experimental) interactions of scientists with these processes or entities. The paper aims to explore how the ordering and structuring of the research narrative relates to the ordering or structuring of the natural narratives.

Historians and sociologists of science, working closely with laboratory notebooks or engaging in ethnographic studies, have pointed out that the accounts of discovery in scientific publications does not fully correspond to the course of events as recorded in the notebooks or observed in the lab. If these texts are carefully fabricated narratives, the question arises what their function is and how it motivates their construction. Research narratives present the procedures that scientists end up with after going through process often involving dead-ends, trial and error, tinkering and calibration, all aspects that are usually omitted from that narrative. Even if the sanitized procedures thus represent results rather than the process, they are still closely connected to the discovery process from which they emerge. Narratives about these established procedures then serve to lend credibility to the data on which interpretations of the phenomena are based, which constitute the second result of research, the natural narrative. In this sense research narratives join the contexts of discovery and justification. Research narratives do so even in a more fundamental manner, by fixing the reference of terms and establishing the relations between entities and their activities that these terms refer to.

The paper thus expands on recent arguments for the important epistemological functions of narratives in science for representation and explanation (Morgan and collaborators, SHPS, Vol. 62) by linking narratives to

questions of experimental practice. To investigate the relation between the research and natural narratives, examples of grafting (or transplantation) experiments in embryology and embryological genetics in the first half of the 20th century are analysed. As has been pointed out, grafting often has an important function in marking processes, such that they can be tracked (Griesemer, 2007). Furthermore, I argue, some of these experiments allow researchers to induce a process and thereby endow the materials in question with the capacity to induce processes. The transfer of agency from the researcher to the material happens in the discursive realm in which the concepts describing the epistemic things involved are formed. Research narratives function to introduce practice in the discursive realm and thereby bring about the transfer.

Inference ‘within the case’: Valency between micro and macro-narratives

Mary Morgan – London School of Economics, United Kingdom

Open field work in the social sciences invariably creates lots of small observational accounts, small because they are slight parcels of life that catch the social scientific gaze, and small because they are treated by the social scientist as the individual jigsaw puzzle pieces that are to be placed into conjunction to make overall sense of the human and social terrain. Analysis of this widespread attachment to micro-narratives suggests the density of this mode of inference, and so raise the question: How does this inference from small narratives to over-all account work?

Two points are important for this characterisation. First, when we look carefully, we can see that, while the individual elements do go to make up the whole, those little units only gain their salience from that broader account just as individual jigsaw puzzle pieces find their meaning only when correctly and closely related with others. And this is not a modular arrangement: there is no simple aggregation just as there is no simple decomposition. Rather, it is a relation of making coherence, of creating valency, between the micro-units, but which also has to work both ways - the macro provides context for the micro, but without the micro there is no content in the macro. Second, the status of these little narratives is hybrid, lying somewhere between data and evidence. Presenting these narratives as chunks of evidence is critical for establishing the credibility and authority of the social scientist as a field worker, both as direct witness to those events, and as a competent observer-cum-analyst in the field. As observant field reports, these little narrative scenes maintain sufficient raw detail to show that the field worker

was present and did observe. But these little narratives also show the field worker's analytical ability to make sense of those observations in ways that can suggest their place of relevance in the overall account, that is - to understand and portray them as 'evidence' rather than data (consistent perhaps with Clifford Geertz's notion of 'thick description'). As a consequence, creating and reporting small narratives out of observational materials means that each chunk is already packaged and presented as a piece of evidence in play that can be used, re-used, perhaps even re-ordered, in making up the broader narrative account of the society.

We see both these features at work in cases as diverse as Merton's study of Mass Persuasion during WWII; the classic 1920s and 1930s studies of US society found in Deep South and in Middletown, and in more recent social anthropology accounts such as Tally's Corner and All Our Kin. While these accounts use different methods of evidence gathering: surveys, interviews, ethnographies and social history - small narratives are characteristic and endemic, found in notebooks through into final reporting. These two features - the creation of coherence and analytical observations - suggest how it is that these little narratives carry valency for inference within the case: from observations, via analysis, to case-findings and macro-narrative. They provide the means, as Geertz suggests 'to generalise within the case'.

Making and Narrating Chemical Beginnings

Matthew Paskins – London School of Economics, United Kingdom

Philosophers and historians of chemistry who are keen to consider the science in its own terms rather than through its relations with physics or biology have sometimes invoked Levi-Strauss' figure of the 'bricoleur': someone who works with available materials and produces patchworks rather than starting from abstract first principles. This has usually been a way of emphasising the tinkering, craft-based aspects of chemical practice. An equally significant problem, however, is how chemists understand what is locally available—the materials and resources with which they begin. This paper argues that the making and narrating of such beginnings offers a useful insight into the ways in which narratives are deployed in scientific practice. Some of the complexities of this process are suggested by a spoof article which was published in *Nature Chemistry* in December 2017, with the title "The Shortest Route to Strychnine." Following a deliberately pompous overview of the "leading concepts of the time" which they had employed, which included "step economy, atom economy, redox economy, word economy, time economy, graduate student economy, and economy", the authors reported that their "efforts were

initiated and concluded by obtaining commercially available strychnine as a light yellow powder from Sigma-Aldrich.” That is, they had bought it from a leading chemical manufacturer.

The paper is a reminder that claims about chemical synthesis are characteristically cast in narrative form. Such claims describe the steps through which a given target can be produced in the laboratory, a process which unfolds in time and by means of specific actions taken by a chemist; they also, self-consciously, refer to previous accounts of how similar achievements have been made before. In addition, the punchline to the article suggests the complex question of where and with what such narratives begin. Total synthesis is supposed to be achieved on the basis of “simple, commercially available precursors”, but both of these factors are shaped by what a chemist happens to have access to. Orienting chemical practices towards these contingencies involves further narration—papers describing organic syntheses will often begin with detailed accounts of the natural history of prized plants, or discussions of industrial conditions which make certain precursors abundantly available as waste products for example.

Such local, improvisatory, beginnings may also be caught up with larger narratives. Such narratives have often been told by chemists in order to justify these new approaches, often with the claim that they systematise practices which are currently disparate and overly particular. Recent examples of such large-scale narratives include E.J. Corey’s retrosynthesis; the claim that combinatorial chemistry can construct libraries of molecules fit for any purpose; and the pervasive story that the introduction of new synthetic materials overturned existing distinctions between nature and artifice. If we grant that such large-scale narratives are crucial in the development of chemistry, this leaves the question of how chemists connect them to their accounts of the resources and materials which are locally available.

Symposium: Joseph Rouse’s philosophy of scientific practice and its implications for the study of the sciences

Organizer: Pieter Present.

Contributors:

Fons Dewulf.

Laura Georgescu.

Pieter Present.

In recent decades, science studies and the history and philosophy of science have been characterised by an increased interest in practice. This “practice

turn” has by now led to well-developed lines of inquiry and fields of study within these disciplines. Increasingly, such studies are taking a self-reflective stance – analysing the notion of “practice” itself, tracing out the implications of the “practice turn” for our understanding of science, and discussing the how to investigate different scientific practices (e.g. Schatzki et al. 2000, Soler et al. 2014).

For the last few decades now, Joseph Rouse has developed a naturalistic understanding of the sciences in terms of research practices that are oriented towards the disclosure and the conceptual articulation of the world. What is especially novel in Rouse’s approach is his philosophical conceptualisation of practices. For Rouse, practices are constituted by the mutual normative accountability of performances. That is, practices are normative, and their normativity is a primitive. What all of this means is that if we are interested in accounting for the sciences as practices, we have to start from the norms of the sciences (in Rousian terms, what matters and what is at stake) and not with the “doings”, or activities, of the sciences. What the activities are, why are they relevant, and why they have the structure they do are determined by the norms involved in the practice(s) to which they belong. With this new take on practices, a wide range of concepts that make up the repository of our current philosophies of science get displaced and reconceptualised, such that the result is very much a specifically “Rousian” outlook on what is at stake in the sciences and in accounts of them.

In this panel, we take this “Rousian” outlook seriously and look at the the implications of Rouse’s philosophy of scientific practices for correlated disciplines, such as the study of scientific education, history and philosophy of science, and the history of philosophy of science.

Schatzki, T.R., K.K. Cetina, and E. von Savigny, eds. 2000. *The Practice Turn in Contemporary Theory*. London: Taylor & Francis.

Soler, Léna, Sjoerd Zwart, Michael Lynch, and Vincent Israel-Jost. 2014. *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science*. New York/London: Routledge.

Rousian History of Philosophy of Science: the Hempel-Neurath debate

Fons Dewulf – Ghent University, Belgium

In this paper I use Joseph Rouse’s conception of scientific practice to discuss the emergence of Carl Hempel’s philosophy of science in the 1940s and 1950s. I argue that, as presented in his works at the end of the 50s, Hempel’s philosophy of science should be strictly seen as a practice centered around the

application of logical models to scientific language in abstraction from the historical, social and institutional aspects of science. The products of Hempel's practice, mostly formal models or definitions, were accountable to intuitions about scientific language, formal toy examples and short summaries of exemplary episodes in the history of science. This meant that Hempel's philosophy of science was only accountable to norms that were accessible to logically trained philosophers. Such a configuration of the practice of philosophy of science installed a discursive regime that could maintain its autonomy not only from the practice of science, but also from the history and sociology of science.

As a case study, I focus on several arguments that Hempel gave in order to defend his covering law model of scientific explanation. I discuss some of his examples to test the formal validity of the model, and the intuitions about explanation that he used to defend his formalization. I show how the first criticisms of the model focused entirely on these aspects alone, and not on scientific activities. However, before Hempel became a central philosopher of science in the US, his practice had not always been like this.

During Hempel's years in Brussels (1934-1939), the norms of his practice were still something at stake. Hans Reichenbach, Rudolf Carnap and Otto Neurath all attempted to change the norms to which Hempel as a philosopher of science should be accountable. As Joseph Rouse has argued extensively in *How Scientific Practices Matter*, Otto Neurath was a representative of a practice-based approach to science. Through my historical research on Hempel's early career, I show that Neurath actively tried to change the aspects of science to which Hempel was accountable in his philosophical practice. I discuss three elements from the Neurath-Hempel correspondence that Rouse picked out as interesting practice based elements in Neurath's ideas. First, I summarize the long debate between Hempel and Neurath over the usefulness of formalization: Neurath constantly emphasized to Hempel the "pragmatic-historicizing" aspect of science over the formal characteristics of scientific language. Second, I discuss Neurath's warnings about the introduction of formal terminology. For Neurath, every introduced piece of terminology would influence the pursuit of science itself. Third, I show how Hempel wanted a philosophical account of truth, while Neurath argued for a use of 'truth' as an element in the scientific deliberation over accepted statements. The Neurath-Hempel debate lasted until Hempel's initial introduction of his model of scientific explanation. Hempel was never convinced by Neurath's pleas for a historical and practical perspective on science. Especially, under the influence of American philosophers in the early 1940s Hempel's formal practice of philosophy of science eventually emerged and became dominant in the Anglophone world.

On how to select entry points in a Rousian history and philosophy of science

Laura Georgescu – University of Groningen, Netherlands

In this talk, I investigate the implications of the Rousian philosophy of scientific practices for how histories of the sciences should be addressed. I take it that the need for such a project goes back to Rouse's conception of scientific practices. For Rouse, practices are, roughly, patterns of interaction between performances that are mutually accountable. This take on practices is opposed to (what Rouse calls) a regularist conception, according to which a practice is an objectively identifiable regular pattern of doings and/or sayings, such that the identification of a practice reduces to the identification of the regularity in question. These divergent philosophical takes on practices have consequences for what it means to provide historical and philosophical accounts of the sciences which begin from and are based on practices. In a regularist framework, scientific practices are what scientists actually do, and historical approaches are reconstructions of such regularities. In a Rousian normative conception of practices, what is at stake in historical accounts of science and how to build such accounts is an open question. This talk takes some steps towards proposing a possible answer.

Thus, I claim that (i) a Rousian philosophy imposes constraints on how to approach the history of science and, (ii) based on Rouse's views about (a) the construction of the temporality of research and (b) how scientific selectivity works, I show what these constraints are. For (a), I discuss the ineliminable prospectiveness of research and show how it is built on a retrospective determination of what is at stake in the research, while for (b) I introduce Rouse's argument for why the sciences are ineliminably selective (and thus cannot be exhaustive) and explain why attending to this selectivity in our philosophical accounts of the sciences is crucial. I also show that, for Rouse, accounts of scientific practices entail making sense of the salience(s) of such practices for whatever research project the account itself is involved in. This applies to both the object of the account and the account itself: that is, it involves reflexivity.

The focus of the talk is on the determination of entry points for reliable historical accounts of science compatible with a Rousian framework. I centre the discussion around the identification of entry points based on (what Rouse calls) "experimental microworlds" (or what Rheinberger calls experimental systems) and on the privileged position that "dead" practices (i.e. practices which have ceased to be reproduced) would have in a Rousian history of science. Experimental microworlds are reliable entry points because they are

points of inflexion in the ongoing “articulation of the world” (which constitute the fundamental norm of scientific understanding), whereas, for “dead” scientific practices, the reconstruction of what was at stake ceases to have immediate salience for ongoing scientific research practices and thus makes them, I argue, true historical products. I will substantiate my philosophical claims with short examples from the early (i.e. pre-nineteenth-century) histories of electricity and magnetism.

Towards a history and philosophy of scientific education in practice

Pieter Present – Vrije Universiteit Brussel, Belgium

Teaching is an important aspect of scientific practice. However, as Lorraine Daston has recently remarked, “we have only the barest beginnings of a history of scientific pedagogy and not even the rudiments of a philosophy” (Daston 2008, 106). Daston refers to some historical studies on scientific education, including the collection of studies edited by David Kaiser (Kaiser 2005). In the concluding chapter, Kaiser and Andrew Warwick provide some general reflections on the usefulness of the works of Thomas Kuhn and Michel Foucault for a philosophy of scientific education (Kaiser and Warwick 2005). Kaiser and Warwick refer to Joseph Rouse, who incorporated insights from both philosophers in his philosophy of scientific practice. However, Kaiser and Warwick merely refer to Rouse as a reader and interpreter of Kuhn and Foucault. In this presentation, I will argue that Rouse’s philosophy of scientific practice deserves to be studied in its own right and that his conceptualisation of scientific practices has important implications for the study of scientific education.

Throughout his work, Joseph Rouse has provided a sustained critique of a reified and representational account of knowledge. In *Engaging Science*, he has further developed his conception of scientific knowledge in terms of practices (Rouse 1996). In this work, Rouse refines his notion of practices, and uses it to present a dynamic view on (scientific) knowledge. Rather than being a thing possessed by knowers, scientific knowledge should rather be seen as a temporally extended process, only existing in and through its continuous repetition and reconfiguration.

In the first part of this presentation, I will argue that Rouse’s dynamic conception of scientific knowledge entails that education should occupy a central place in our analyses of scientific practices, as it is crucial in guaranteeing their temporal extension and sustenance. Starting from the “ten theses on practice” which form the backbone of *Engaging Science*, I show how each of these theses leads us to a consideration of scientific education.

However, Rouse's reconceptualization of scientific knowledge also has implications for our understanding of scientific education itself. In the second part of this presentation, I will work out these implications, focussing on Rouse's non-subject-centered account of scientific practices. Rouse insists that practices "are not just agents' activities but also the configuration of the world within which those activities are significant" (Rouse 1996, 133). In line with this, I will argue that we should analyse scientific education as not (only) an individual process, but a distributed one. This process includes the continuous reconfiguration of institutions, teaching materials and spaces, among other factors.

Daston, Lorraine. 2008. "On Scientific Observation." *Isis* 99 (1):97–110.
Kaiser, David, ed. 2005. *Pedagogy and the Practice of Science: Historical and Contemporary Perspectives*. Cambridge, MA: MIT Press.
Kaiser, David, and Andrew Warwick, 2005, "Kuhn, Foucault, and the Power of Pedagogy" in (Kaiser 2005), pp. 393-409.
Rouse, Joseph. 1996. *Engaging Science: How to Understand Its Practices Philosophically*. Ithaca/London: Cornell University Press.

Symposium: Philosophy of Science Policy & Pragmatism

Organizers: Jorrit Smit, Erman Sözüdoğru.

Contributors:

Chiara Ambrosio.

Stephanie Meirmans.

Jorrit Smit.

Erman Sözüdoğru.

This symposium brings together a broad range of philosophers interested in the less explored nexus of science policy and pragmatism. The focus is partly on the methodological, epistemic, and ethical questions that rise from understanding scientific knowledge in policy contexts as practices. Also, the session will display a variety of evidence that demonstrates that scientific practices and knowledge are shaped by their broader context in particular policies. We aim to investigate what we could gain (or lose) from thinking in pragmatist terms about the principles and practices of science policy, moving from the 'economy' of research in classical pragmatism to the 'ethics' of research in current practice. Our discussion follows on from original pragmatist texts as well as contemporary literature on values and the role of science in policy making (Douglas 2009, 2013; Brown 2012, 2013; Kitcher 2003, 2016) Our

aim is to contribute to the existing literature on pragmatism and science policy by combining conceptual approaches to science policy with considerations of its current problems in case studies, all informed by or taking issue with pragmatism.

The first two papers will address the thought of classical pragmatists like Charles Peirce and John Dewey, and assess their philosophical value for the issues we face today. For example, can Peirce's 'Economy of Research' be helpful in overcoming the divide between 'useless' theoretical science and matters of practical utility? The concept of utility is central to science policy and will be critically scrutinized in dialogue with pragmatist philosophies. This will allow altered perspectives on current science policy concepts like valorisation and impact: the explicit demand that science should become societally valuable.

The other two papers will take current problems in the governance of scientific research as their starting point to reflect on the potential merits of a pragmatist approach. In policy efforts to eliminate specific diseases, for example, a plurality of scientific practices requires consideration. It is argued that policy could benefit from a pragmatic framework to assess the variety of values at play in determining the limits of plurality. The other example from practice will consider the adverse effects of competitive research funding. How might we be able to apply our (pragmatist) insights regarding science policy to this issue, and could it open up a future of a more responsible practice of research?

The main aim of this symposium is to foster an exchange of ideas between scholars working on topics regarding science policy and pragmatism, and to discuss ideas and tools for how to develop a (pragmatist) philosophy of science policy.

The session will consist of relatively short interventions by the four presenters, which will leave ample time for discussion. We will actively invite the participation from the audience as well.

“Running a Steam Engine by Burning Diamonds”? Values and Virtues in Peirce’s Economy of Research

Chiara Ambrosio – UCL, United Kingdom

In her study on “The Moral Economy of Science” (1995), Lorraine Daston includes Charles S. Peirce in a broader discussion of the ethos of coordination and international collaboration that is distinctive of late nineteenth century science. Daston’s reference to Peirce is part of a broader claim, aimed at showing the value-ladenness of a set of practices distinctively associated to

certain forms of empiricism – particularly quantification and objectivity. But Peirce also wrote specifically, and directly, on the material economy of research, which he saw as an important constraint on the pursuit of hypotheses (McKaughan 2008, Nyrup 2015). This started as a very practical problem, which is now recognised as one of the earliest cost-benefit analyses of scientific research (Wible 2006). In 1879, while working on gravimetric measurements for the US Geodetic and Coast Survey, Peirce produced a set of “Notes on the Economy of Research” (W4, 72-78) exploring “the relation between the exactitude of a result and the cost of attaining it” (ibid. 73). Incidentally, the notes immediately followed the publication of his key pragmatist writings, “The Fixation of Belief” and “How to Make our Ideas Clear” (1878). Peirce carried on elaborating on economic considerations, but in subsequent years his attention shifted from measurement and exactitude of results to the economy of research as an evaluation of the cost and value of a project, as well as its possible effects on other projects (EP2, 107ff). At the same time, he developed an apparently idiosyncratic approach to theoretical science, which he characterised as “the study of useless things. For the useful things will get studied without the aid of scientific men. To employ these rare minds on such works is like running a steam engine by burning diamonds” (CP 1.76). In this contribution, I will explore this apparent idiosyncrasy and reconcile these two opposing strands – theoretical science vs. matters of practical utility – under Peirce’s distinctive brand of pragmatism. In doing this, I will depart from traditional interpretations of the economy of research as a by-product of formal debates around theory-choice, and show instead the ways in which it stands out as an early chapter in the integrated history and philosophy of science policy. I will conclude my intervention showing that Peirce’s pragmatism reconciles, historically as well as conceptually, material and moral aspects of the economy of science.

Funding policies: How to foster responsible research practices?

Stephanie Meirmans – University of Amsterdam, Netherlands

Competitive research funding has an immense impact on science. Ideally, funding agencies should be able to select the best researchers and scientific projects and provide them with sufficient resources. However, in recent years it is becoming more and more obvious that funding systems cannot live up to these expectations; that they indeed may even have adverse effects. It has been suggested that the system may lead to scientists chasing their h-indexes rather than performing responsible research, to researchers spending more time on writing funding applications than doing research or supervising PhD

students, and to scientists cutting corners to be the first instead of being right. One main underlying reason for such trends is the current practice of research evaluation in quantitative rather than qualitative terms (i.e. counting the number of publications rather than their scientific value), and voices get louder demanding a system that enhances quality over quantity.

In this talk, I first argue that we should improve the system by having funding policies that foster responsible research practices. Importantly, this includes a refocus from evaluating products to fostering processes of doing responsible science. I will then address what might be such responsible research practices by using insights from scientific practice in evolutionary biology and by integrating those with insights from philosophy, gender studies and STS. This symposium provides a unique chance to compare these insights to pragmatist thinking, which seems to at least partly overlap.

Utility & Pragmatism: situating science policy's purpose

Jorrit Smit – Leiden University, Netherlands

Utility in a limited economic sense has dominated the legitimation and practices of public funding of scientific research in the post-war Western world. The 'creation of value from knowledge by making it appropriate or available for societal and economic use' is still today the purpose of science policy, captured in the Dutch policy concept of 'valorisation'. This relatively novel science policy concept has by now materialized into specific valorisation paragraphs in grant proposal assessments at funding organizations, valorisation support and education at universities, as well as rankings and indicators to evaluate 'valorisability' of research, researchers, and universities. In this paper, I develop from a pragmatist perspective an understanding of utility of scientific research that goes beyond mere results, products, or outcomes. Informed by the philosophies of John Dewey, Joseph Rouse, and Isabelle Stengers, scientific research is situated as a social, material and discursive practice in a democratic society. This places emphasis on its direct and indirect consequences as well as the concerns of a public and the care of political forms that emerge in response. Concepts like economic utility and valorisation re-appear as specific manifestations of this situated value of scientific practices.

A pragmatic reading of scientific research poses questions to past and current science policies and funding practices: how did different understandings and manifestations of utility of science played a part in the process of planning science for the public good – and what consequences for the practice of research did this have? The philosophical understanding of utility proposed

here, should help understand different manifestations of utility as historical instantiations of this more general aspect of research as practice. Ultimately, this invites reflections on the possibility of pragmatic interventions in science policy: could it make an alternative valuable science possible that takes value-laden research as starting point?

Pragmatic Pluralism: investigating the boundaries of plurality in neglected tropical disease research

Erman Sözüdoğru – UCL, United Kingdom

In this paper I aim to provide a pragmatic framework to investigate limits of both epistemic and methodological pluralities in scientific practices. In particular, I will focus on the current efforts directed by the World Health Organisation to eliminate Human African Trypanosomiasis (HAT, also known as the sleeping sickness) by 2030. Using this case study I will argue that the plurality of approaches in a given field of scientific practices (in this case aiming to eliminate HAT) are limited by a broad range of socio-economic and political values. Here I develop a pragmatic framework to critically scrutinise these values in science policy context which pose boundaries on the plurality of approaches and accounts used in scientific practices.

I take pluralism to be the rejection of the monist assumption that the aim of science is to provide a single, complete and coherent account of phenomena. Instead, I argue that monist assumptions must be challenged and replaced with the following pluralist tenets: there are multiple aims in science; different approaches have distinct aims, focusing on different aspects of phenomena; and each account is particular to the specific questions and aims of an approach.

Herein I focus on current efforts of the World Health Organisation to eliminate HAT – in particular, the development of new anti-parasitic drugs. I argue that drug discovery and development requires a plurality of approaches, each focusing on different aspects of phenomena. The pluralist argument I support here is normative in the sense that scientific inquiry ought to be pluralist (instead of monistic), in which a multiplicity of accounts and approaches is necessary to explain and explore different aspects of phenomena. Moreover, I argue that the plurality of approaches and accounts employed to achieve a certain aim is bounded by pragmatic values. I argue that pragmatic values determine the best way to achieve a specific aim within the broader socio-economic and political context of scientific inquiry. Here, I argue that the extent of plurality in scientific practices involved in developing new drugs to eliminate HAT must be understood with respect to the pragmatic values that

define the best way to eliminate HAT in its current socio-economic and political context.

In this paper, I support a normative argument for pluralism, challenging monist assumptions about scientific practices and their aims. Moreover, I develop a pragmatic framework within which to understand and explain the extent of pluralities in scientific practices. In short, I take the motto of pluralism to be 'many things go'. Where the pragmatic framework I aim to develop here allows for a close investigation aims and values of scientific practices to answer the question what goes and why? I argue that asking this question will allow for critical scrutiny of values in science policy context, which pose limits on the plurality in scientific practices.

Symposium: Values in the Practice of Archaeology and Paleontology

Organizers: Derek Turner, Rune Nyrup, Caitlin Wylie, Joyce Havstad.

Contributors:

Rune Nyrup.

Derek Turner.

Caitlin Wylie.

Much of the recent work in the philosophy of historical science (e.g. paleontology and archaeology) has focused narrowly on epistemological issues: How do scientists reconstruct the deep past? What are some of the limitations and obstacles that they confront? And what accounts for their epistemic successes? But some other important aspects of the practice of paleontology and archaeology have received relatively less attention. What role, for example, do non-epistemic (including aesthetic, ethical, social, and/or political) values play in paleontological and archaeological practice?

Beginning with Heather Douglas's (2002) argument from inductive risk, many philosophers interested in understanding the role of non-epistemic values in science have tended to look at cases involving policy-relevant science, such as biomedical or climate science. Paleontology and archaeology are an especially interesting place to analyze the role(s) of non-epistemic values, precisely because efforts to reconstruct the past are less directly related to policy. How might the issue of inductive risk come into play in paleontology and archaeology? Are there other ways that non-epistemic values might figure internally in the practice of archaeology and paleontology? If thinking about non-epistemic values can help to illuminate the practice of paleontology and archaeology, that might suggest that various non-epistemic values will turn out

to be doing important work in surprising places. The papers in this session seek in various ways to broaden our understanding of the practice of paleontology and archaeology by moving beyond the narrowly epistemic focus that has characterized work in this area over the last fifteen years.

How Archaeologists Resolve the Inductive Risk Argument

Rune Nyrup – University of Cambridge, United Kingdom

Feminist archaeologists (e.g. Gero 2007, *Journal of Archaeological Method and Theory*) have recently criticised the tendency to value certainty and avoid ambiguity in archaeological interpretation. This, they argue, has led archaeologists to primarily tell stories about the parts of past societies where materially rich and powerful people left their mark. This is not a value-neutral decision, as it systematically disincentivises archaeological work on the lives of non-elite populations (e.g. women, servants, rural populations) which are already underrepresented in historical narratives.

In effect, this is a version of the inductive risk argument familiar from discussions of values in policy-relevant sciences: in setting high standards for accepting interpretations, traditional archaeology has put too much emphasis on avoiding false positives (accepting false or unsupported interpretations) and too little emphasis on avoiding false negatives (the many underexplored topics).

However, unlike the typical inductive risk arguments concerning policy-advising sciences, the upshot for proponents of this argument in archaeology is not that archaeologists should re-adjust the evidential threshold for accepting an interpretation as true. Rather, Gero argues that archaeologists should work on making the uncertainties and ambiguities involved in archaeological interpretation more explicit and give more professional recognition to work that highlights and discusses these, without necessarily lowering the evidential standards of the field. This strategy is closer to how Betz (2013, *European Journal of Philosophy of Science*) defends the value-free ideal against the inductive risk argument.

This paper analyses what the relevant differences for the inductive risk argument are between archaeology and policy-advising sciences. I argue that the crucial difference concerns the route by which the conclusions drawn by scientists can have value-laden consequences. In policy-advising sciences, this is through the actions of policy-makers. It arises from the policies that will be implemented (or not) based on whether scientists accept a given hypothesis, say, concerning the safety of some chemical. In archaeology, by contrast, the relevant consequences include representational harms arising from whether

and how different groups are represented within archaeological narratives. Thus, whereas the potential goods and harms of policy-advising sciences usually stem from actions taken on the basis of accepted hypotheses, the inductive risks highlighted by feminist archaeologists are, at least in part, more directly contained in the conclusions archaeologists accept, reject or neglect. If representational harms were the only relevant concern in archaeology, this would give rise to a variant of the problem of wishful thinking: it might be used to argue that archaeologists should simply invent the narratives that have the best representational effects, with no regards to the evidence. To avoid this problem, I outline an account of the value of having evidentially well-supported knowledge of past societies. On the basis of this account, I argue that Gero provides a plausible resolution of the inductive risk argument in archaeology, by promoting methodological standards which strike a balance between pursuing the value of evidentially supported knowledge and avoiding representational harms.

The Functional Beauty of Fossils: Aesthetic Values in Paleontological Reconstruction

Derek Turner – Connecticut College, United States

Paleontologists' efforts to infer the functions of fossilized structures have an important aesthetic dimension. In their 2008 book, *Functional Beauty*, Glenn Parsons and Allen Carlson, researchers best known for their contributions to environmental aesthetics, elaborate and defend the claim that an item's positive aesthetic qualities often have something to do with its fitness for function. They draw upon a version of the selective-historical account of biological functions in order to undergird this theory. Parsons and Carlson are committed to cognitivism in aesthetics, which is the view that appropriate aesthetic appreciation of an object requires knowledge about it. On this view, knowing an item's biological function deepens and enhances our aesthetic appreciation of it.

Parsons's and Carlson's view has unexplored implications for our understanding of the practice of paleontology. In particular, their view suggests that the inference of function from fossilized structure may serve aesthetic as well as epistemic goals, and that research in functional morphology is a form of aesthetic engagement with fossils. One example of a pattern that poses this sort of inferential problem is the trend toward greater ornamentation—more ribs, nodules, and spines—in the shells of Mesozoic ammonoids. What function, if any, did the ornamentation have? That this is partly an aesthetic issue can be brought home by looking at how scientists are

beginning to tackle it. One new approach is to use 3D printing technology—already widely in use by artists—to create plastic models of ammonoid shells with different geometries. Those models can then be used in crushing experiments to see whether the ribs, nodules, and spines made shells more resistant to crushing by predators. This experimental practice is a way of probing the functional beauty of ammonoid shells.

Assumptions about Time in Specimen Research

Caitlin Wylie – University of Virginia, United States

To become physically stable and epistemically reliable enough to study, artifacts and specimens undergo significant processing. Once collected, these objects are cleaned, repaired/reassembled, sometimes coated with glue, and given a label, a storage container, and a database entry. This work crucially shapes how an object looks and thus how a scientist interprets it. Based on interviews and observations of workers in fossil laboratories, I found that practitioners – i.e., scientists, preparators, and conservators – hold divergent beliefs about time with regards to specimen processing, and that these beliefs inform how they shape and study fossils.

These values encourage practices that range from patient, long-term planning (inspired by fossils' incomprehensibly old age and by practitioners' desire to do the best possible work) to frantic urgency (driven by scientists' need to publish to secure their jobs and their intellectual claim to new knowledge). Specifically, conservators follow an often-repeated goal of preserving objects for the next 100 years, which guides their choices of chemicals (i.e., ones that don't degrade) and interventions (i.e., minimal, to reduce risk of introducing new sources of degradation). They perceive objects as valuable in themselves and therefore worth preserving. In comparison, scientists understand fossils as a means to an end, i.e., publications. They believe that the papers and intellectual property endure, not necessarily the fragile fossils. On the other hand, most publications include analyses of already-prepared fossils in collections; thus conserving fossils promises future research objects and opportunities. Preparators have traditionally followed scientists' short deadlines, which meant using instant glues (which set quickly but degrade in a few decades) and removing all the rock around fossils (making them fully visible but physically weak). More preparators have begun to adopt conservators' values, such as by choosing chemically-stable adhesives, even though they set more slowly, and by removing rock only from the scientifically important parts of fossils, which improves work efficiency and leaves more physical support for the fossils. The conflicts within and between groups about

what makes for “good” specimens shapes how they work with and understand fossils. As a result, practitioners’ notions of time shape their and our conceptions of the organisms and environments of deep geological time.

Contributed Talks (alphabetical by last name of lead author)

Representation Re-construed: Answering the Job Description Challenge with a Construal-based Notion of Natural Representation

Mikio Akagi – Texas Christian University, United States

William Ramsey (2007) articulates a common complaint that cognitive scientists apply the concept “representation” too liberally. He argues that representations are often ascribed according to a crude causal theory he calls the “receptor notion,” According to which a state s represents a state of affairs p if s is regularly and reliably caused by p (119). Ramsey claims that the receptor notion is what justifies the ascription of representations to edge-detecting cells in V1, fly-detecting cells in frog cortex, and prey-detecting mechanism in Venus flytraps. However, Ramsey also argues that the receptor notion justifies ascribing representational states to the firing pin in a gun: since the state of the trigger regularly and reliably causes changes in the state of the firing pin, the firing pin represents whether the trigger is depressed. The firing pin case is an absurd consequence. He concludes the receptor notion is too liberal to be useful to scientists.

I argue that something like Ramsey’s receptor notion can be salvaged if being a receptor is contextualized in terms of “construal.” Construals are judgment-like attitudes whose truth-values can vary licitly independently of the situation they describe. We can construe ambiguous figures like the Necker cube as if it were viewed from above or below, and we can construe the duck-rabbit as if it were an image of a duck or of a rabbit. We can construe an action like skydiving as brave or as foolhardy, depending on which features of skydiving we attend to. On a construal-based account of conceptual norms, a concept (e.g. “representation”) is ascribed relative to a construal of a situation.

I describe a minimal sense of what it means to construe a system as an “organism,” and how ascriptions of representational content are made relative to such construals. Briefly, construing something as an organism entails construing it as having goals and mechanisms for achieving those goals. For example, frogs-qua-organisms have goals like identifying and ingesting food. I suggest that ascriptions of natural representations and their contents are always relative to some construal of the representing system qua organism.

Furthermore, the plausibility of representation-ascriptions is constrained by the plausibility of their coordinate construal-qua-organism. So the contents we

ascribe to representations in frog visual cortex are constrained by the goals we attribute to frogs.

Absurd cases like Ramsey's firing pin can be excluded (mostly) since guns are not easily construed as "organisms." They do not exhibit autonomous behavioral dynamics, and they are difficult to anthropomorphize. It is not impossible to ascribe goals to weapons or other tools, but the ascription of folk-psychological properties to tools generally follows a distinct pattern from representation-ascription in science.

My construal-based proposal explains the practice of representation-ascription better than Ramsey's receptor notion. It preserves Ramsey's positive examples, such as the ascription of representations to visual cortex, but tends to exclude absurd cases like the firing pin. Since cognitive scientists do not actually ascribe natural representations to firearms, I submit that my account is a more charitable interpretation of existing scientific practice.

Simulation, Test bench, and Hardware-in-the-loop: Validation in engineering design processes – A case study

Sabine Ammon – TU Berlin, Germany

Henning Meyer – TU Berlin, Germany

Whereas the primary aim of science is to generate new knowledge, the primary aim of design is to create new artefacts (such as products, processes, or systems). Just as scientific methods are employed to make the endeavor of science successful by following a rigidly reviewed, reliable, and systematic course of action, design methods are used to make the endeavor of design successful by following a reliable and systematic course of action (while review and the requirement of systematicity usually are less pronounced). At second glance, however, design is as much related to the generation of new knowledge as it is science, albeit more intrinsically: the design process is essentially about understanding the future artefact; hence, designing itself is an epistemic activity, which results in a knowledge about the future artefact.

In order to learn more about the knowledge dynamics in engineering design processes, our contribution examines the role of simulation models and their respective validation strategies by drawing on a case study. The goal of the InDriVe project was to design a gear for a vehicle simulator, which is able to simulate, test and experience innovative vehicle and drive concepts. As a tool, InDriVe Hybrid Simulator can be used in early phases of the development of future cars. By using this example, we investigate how modelling practices allow for a better understanding of the future artefact, how validation takes place in these different phases of the product development and to what

extend these procedures contribute to a gain of knowledge.

Drawing on the example of the InDrive-Hybridsimulator we find three different instances of testing: a FEM-model during the construction phase, the prototype on the test bench, and the engine and drive train in the vehicle simulator. The FEM-model (based on the finite element method resp. analysis) allows to simulate stress in the component and to perform a stress analysis. The prototype on the test bench allows to simulate the vehicle operation under idealised standard condition in order to test specific parameters. The vehicle simulator couples computer-based, simulated feedback control systems with “hardware-in-the-loop” methods in order to induce different kinds of longitudinal dynamics. Their impact allows to predict the driving performance of the future vehicle; insights which, in turn, are used by the car company to decide on a new line of products.

As epistemic tools (Boon & Knuuttila 2009) those simulation models and their validation strategies are embedded in complex milieus of reflection which comprise soft- and hardware constellations, operational procedures, notations, artefacts as well as the expertise of the developer. With the help of our analysis we want to shed light on these different developmental milieus in action, which not only allow to create a new product (the gear, the vehicle), but also to generate knowledge about the future artefact.

Boon, Mieke and Tarja Knuuttila, Models as epistemic tools in engineering sciences, in: Anthonie Meijers (Ed.), Philosophy of technology and engineering sciences, Handbook of the philosophy of science, vol. 9, Amsterdam 2009, pp. 693-726.

Philosophy of scientific malpractice

Hanne Andersen – University of Copenhagen, Denmark

Misconduct and questionable research practices (QRP) are major concerns in contemporary science. Empirical studies indicate that sloppy or negligent practices are more prevalent than outright misconduct [1-3], and that scientists worry less about purposefully manipulated data than about the grey zone between ‘cleaning’ and ‘cooking’ data, or about how decisions regarding sorting, interpretation and inference should be made and reported [4]. Hence, in analyzing malpractice an important task is to distinguish how neglect or recklessness differs from intentional fraud. To contribute to a nuanced understanding of how scientific misconduct and QRP can be detected and corrected as early as possible in the research process, this paper outlines an epistemological analysis of recklessness and negligence in science.

Comparisons will be made to existing analyzes of which categories of QRP lie behind retractions and corrections in the journal literature [7-11]. On this basis, the talk will conclude with suggesting a new approach to responsible conduct of research (RCR) training that provides scientists in their roles as collaborators and peers with the tools to intervene at an early stage.

1. Martinson, B.C., M.S. Anderson, and R. De Vries, Scientists behaving badly. *Nature*, 2005. 435: p. 737-738.
2. Fanelli, D., How many scientists fabricate and falsify research? A systematic review and meta-analysis of survey data. *PLoS ONE*, 2009. 4(5): p. 1-11.
3. Bouter, L.M., et al., Ranking major and minor research misbehaviors: results from a survey among participants of four World Conferences on Research Integrity. *Research Integrity and Peer Review*, 2016. 1(1): p. 1-17.
4. De Vries, R., M. Anderson, and B.C. Martinson, Normal Misbehavior: Scientists Talk About the Ethics of Research. *J Empir Res Hum Res Ethics*, 2006. 1(1): p. 43-50.
5. Allchin, D., Error Types. *Perspectives on Science*, 2001. 9: p. 38-59.
6. Hon, G., J. Schickore, and F. Steinle, Going Amisss in Experimental Research. *Boston Studies in the Philosophy of Science*. Vol. 267. 2009, Berlin: Springer.
7. Nath, S.B., S.C. Marcus, and B.G. Druss, Retractions in the research literature: misconduct or mistakes? *MJA*, 2006. 185(3): p. 152-154.
8. Redman, B.K., H.N. Yarandi, and J.F. Merz, Empirical developments in retraction. *Journal of medical ethics*, 2008. 34: p. 807-809.
9. Wager, E. and P. Williams, Why and how do journals retract articles? An analysis of Medline retractions 1988-2008. *Journal of medical ethics*, 2011. 37: p. 567-570.
10. Steen, R.G., Retractions in the scientific literature: do authors deliberately commit research fraud? *Journal of medical ethics*, 2011. 37(2): p. 113-117.
11. Hosseini, M., et al., Doing the right thing: A qualitative investigation of retractions due to unintentional error. *Science and Engineering Ethics*, 2017.

The epistemological significance of repertoires: Tools to understand representational attributions

Rachel Ankeny – University of Adelaide, Australia

Sabina Leonelli – University of Exeter, United Kingdom

It is widely acknowledged that models come in an endless variety of forms, a combination of which is always required by their use in scientific practice. Given this dramatic diversity, much attention has been paid to the actual features of models employed in scientific practice in order to clarify the

epistemological status of each type of model as both a product of and a tool used for scientific theorizing (e.g., Weisberg 2013; Levy and Currie 2015; Frigg and Nguyen 2016). Relatively less attention has been devoted to the variety of activities, such as extrapolation, that need to be performed to yield models that can be defined as ‘good’ or ‘adequate’ (cf. Steel 2008; Knuuttila 2011; Baetu 2016).

Examining modelling activities, rather than their products, is a particularly useful approach when trying to understand how experimental organisms help to create knowledge that can be projected beyond the immediate domain in which it was produced, and particularly what makes such projections more (or less) plausible. This question is especially significant given that organisms often are taken as models for phenomena that are arguably not directly observable in the organisms themselves (e.g., the use of mice to explore alcoholism in humans) or for organisms that are very dissimilar to them (e.g., the use of yeast as models for cancer in humans).

In this paper, we argue that the plausibility of organisms as models relates to the ways in which they fit (or fail to fit) a given research repertoire, which in turn defines the expectations and constraints of the research community in question. We thus provide a philosophical framework to understand the epistemic grounds on which researchers endow models with representational power, the extent to which such endowment is viewed as fruitful and plausible—or problematic and unrealistic—by others, and the implications of such assessments for what is perceived as ‘successful’ research practice. This analysis also illustrates one way in which adopting the framework of repertoires can help address long-standing questions within the philosophy of science.

Structuralism and the Metaphysics of Biological Practice

Tiernan Armstrong-Ingram – University of California, United States

Structuralism purportedly resolves problems of objecthood in mathematics, physics, and scientific theory change. Steven French has argued that the ‘structuralist tendency’ in philosophy of science ought to be extended in the form of Biological Structural Realism (BSR). He argues that the success of Ontic Structural Realism (OSR) for fundamental physics suggests that BSR will produce similar successes for biological problems. There are two reasons to be skeptical of French’s move from OSR to BSR. First, strong disanalogies exist between the problem of individuality in fundamental physics and the various problems of biological individuals. The disanalogies suggests that structural solutions to biological problems, if they exist, may not bear any strong

resemblance to OSR. Second, OSR is developed as a metaphysics **for** physics, a response to specific problems and practices within the science of physics. French develops his BSR primarily as an extension of OSR and not in response to problems faced within the actual practice of biology. BSR is not developed as a metaphysics **for** biology. Contra French's suggestion, structuralism is already present in the metaphysics of biological practice and intertwined with the problems of biological individuality. The range and diversity of problems and practices in biology suggest that a monist metaphysics will be inadequate for all biology. The more plausible alternative to BSR is to explore and develop a pluralist metaphysics for biology, inclusive of objects, processes, and structures.

Virtual Morris Water Maze: The Independent Life of an Experimental System

Nina Atanasova – The University of Toledo, United States

The premise of this paper is that advancements in contemporary neuroscience are best represented by developments of new experimental techniques and therapeutic innovations afforded by novel technologies. Neuroscience, similarly to other biomedical sciences, has increasingly become data-driven. These observations motivate the view developed in this paper, namely that contemporary neuroscience progresses through technological innovations as opposed to guidance of overarching theories.

This view challenges Kuhnian accounts of science such as the ones developed in Longino (2013) and Sullivan (2017). Even though the two differ in their accounts of scientific progress, both views share the Kuhnian assumptions that science is theory-driven, observation is theory-laden, and different scientific communities are bound by different theoretical, hence ontological, commitments. According to both accounts, the experimental results and their corresponding explanations produced in different scientific communities are thus incommensurable. For Longino, maintaining multiple overlapping though incommensurable approaches is, in fact, conducive to scientific progress. In her view, the challenges raised by different epistemic communities against each other, force all of them to improve and strengthen their individual knowledge producing practices. For, Sullivan, on the other hand, progress ideally consists in theoretical unification of the otherwise incommensurable approaches. She, thus, advocates unification of scientific vocabulary and standardization of experimental techniques with the goal of integration of neuroscience.

The alternative proposed here is inspired by approaches to the study of science in practice advocated by Rheinberger (1992) and Ankeny and Leonelli (2016)

among others. Shifting the focus of the analysis of scientific progress from theory development to social and material practices of knowledge production, shows that while scientific communities are diverse and don't always share ontological commitments, they can nevertheless cooperate and produce integrated accounts of the phenomena of common interest. However, the progress of neuroscience cannot be reduced to the improvement of theoretical claims produced in this process. Rather, progress is measured with the generation of new hypotheses and experimental techniques articulated in this process of scientific cooperation and collaboration aided by new technological advancements.

The case for this view is illustrated by the continuous modifications of one of the most widely utilized behavioral tests in neurobiology, the Morris Water Maze. It is an experimental system initially developed as a test for rat learning and spatial memory (Morris 1981). The system has been subsequently adapted for mice and most recently for humans. In the case of humans, it is used as a virtual navigation task in neuroimaging studies. The latter development was made possible with the advancement of functional magnetic resonance imaging.

I argue that it is the opportunism afforded by the developments of new technologies that often guides the experimental process in science. In the Morris Water Maze case, the availability of neuroimaging technology enables neuroscientists to perform experiments on human subjects modeled after previously successful experiments with rodents. In this case, technology rather than some overarching theory shapes the experimental and knowledge generating practices in the field.

Ankeny, R. A. and S. Leonelli (2016). "Repertoires: A post-Kuhnian perspective on scientific change and collaborative research". *Studies in History and Philosophy of Science* 60: 18-28.

Longino, H. E. (2013). *Studying Human Behavior: How Scientists Investigate Aggression and Sexuality*. Chicago: University of Chicago Press.

Morris, R. G. M. (1981). "Spatial Localization Does Not Require the Presence of Local Cues". *Learning and Motivation*, 12: 239-260.

Sullivan, J. A. (2017). "Coordinated pluralism as a means to facilitate integrative taxonomies of cognition". *Philosophical Explorations*, 20(2): 129-145.

Rheinberger, H. (1992). "Experiment, difference, and writing: I. Tracing protein synthesis". *Studies in History and Philosophy of Science*, 23(2): 305-331.

Visual representations as evidence in cell biological mechanism research

Yin Chung Au – National Cheng Kung University, Taiwan

This study explores the practice of using data images for constructing cell biological mechanisms. It seeks to answer the following questions. How are visual representations of cell events and cellular/ biochemical entities translated into evidence for mechanisms? How does researcher define 'evidence' at different stages of reasoning? The distinction between evidence and representation is sometimes vague in practitioner's language. In scientific papers, the nomenclature of visual evidence is normally technique-based, such as 'histological evidence' and 'immunostaining evidence'. These names nonetheless suggest more the means of producing representations than the evidence for specific arguments. How are evidence 'kinds' actually defined in the practice of searching cell mechanisms?

Cell biology research is mechanism-oriented, where visualisations play a crucial role in determining causal relations of entities and activities, which are the components of cell mechanisms (Craver and Darden 2013). Some literature has elaborated the verbal-visual interaction in the construction of scientific arguments (Gross and Harmon 2014). Amann and Knorr-Cetina's pioneering study (1988) shows how visual data is consensually fixed as evidence via laboratory conversations. Yet the fine distinction between representation and evidence in the practice remains less-explored. Also, the proliferation of imaging technologies has resulted in a great variety of new kinds of visual representation. The arising question is whether the 'kinds' of techniques, representations and evidence are consistently defined in the practice.

This study proposes a hypothetical model of the practice of inter-translating visual representations and visual evidence. First, the researcher decides the representational meaning of data. Second, the representation of event/ entity becomes the evidence for a specific mechanism component. Third, this evidence is in turn used as a representation for surrogative reasoning in the process of organising components of the pathway concerned.

In this process, different kinds of representation may be treated hierarchically according to their usefulness for determining causal relevance and causal productivity (in Stuart Glennan's term). Both kinds of causal relation need to be determined for constructing a complex mechanism. Normally, researchers first determine causal relevance and then causal productivity. The more likely can causal productivity be inferred from the representation, the greater validity does the representation acquire as evidence. However, a mechanism with defined boundary can still contain lots of confounding variables that are

unknown and/or partially known. Due to such complexity, while some ‘direct’ kinds of visual representations are viewed as ‘decisive evidence’, they alone are not sufficient for constructing mechanisms if causal relevance has not been previously inferred from less direct representations.

This study combines a qualitative analysis of scientific papers and a laboratory study. Papers analysed were randomly sampled from a previously established database of papers in the apoptosis (programmed cell death) field.

Navigating Evidential Discord in Observational Cosmology

Michael Begun – University of Pittsburgh, United States

One intriguing feature of recent research on some prominent questions in astrophysics and cosmology is the presence of stubborn discrepancies or tensions in empirical results. For example, the two main approaches for determining the Hubble constant—classical determinations involving distance ladders typically calibrated by Cepheids and Type Ia Supernovae, and CMB-based determinations—have tended to yield discrepant results, with the most precise recent estimates of each type producing a 3.4σ difference. More controversially, in the context of dark matter research, the research group behind the DAMA/LIBRA experiment has claimed to find a strong signal for the existence of WIMP dark matter, whereas others have claimed to rule out the existence of dark matter in the same parameter range, yet no obvious explanation for the discrepancy has emerged. While evidential discord is by no means unique to astrophysics and cosmology, cases like these provide a useful starting point for understanding evidential discord more generally, as well as highlighting some of the unique empirical challenges facing astrophysics and cosmology.

In this paper, I examine the discrepant results in the Hubble constant and dark matter cases and use them to try to better understand the ways in which empirical results can conflict and the epistemic implications of those conflicts. Starting from Jacob Stegenga’s account of inconsistency and incongruence, I argue that a more nuanced picture of evidential discord is required for making sense of the Hubble constant and dark matter cases. I characterize evidential “non-conformity” as a weaker form of discord than inconsistency but a stronger form than incongruence, and show that the Hubble constant and dark matter cases fit this characterization of non-conformity. One reason why the results in these cases are better characterized as non-conforming rather than inconsistent is that because the competing approaches rely on different methodologies and background assumptions, the discrepancies may ultimately be found to be compatible, perhaps through the modification of background

assumptions or with the discovery of currently unknown physical features affecting the results. I also show why the evidential discord in the Hubble constant and dark matter cases should not be characterized as incongruent on Stegenga's definition.

In the final part of the paper, I examine the current prospects and scientific strategies for resolving the discrepancies in the Hubble constant measurements and in the dark matter detection experiments. There is now a strong push for new, more precise measurements and experiments, reexaminations of experimental methods to uncover systematic errors, and critical inspections of physical assumptions. I suggest that whereas judgments of evidential non-conformity are likely to be experimentally fruitful, leading to improved experiments and methodologies, judgments of inconsistency are more likely to be theoretically fruitful, leading to revised models or theories. While evidential discord is often seen by philosophers of science as a serious problem, this analysis highlights the positive epistemic role that it plays, at least in contemporary astrophysics and cosmology.

Similarity 'All the Way Down': A Cognitive Approach to Pluralistic Realism

Corinne Bloch-Mullins – Marquette University, United States

Pluralistic realism maintains that, while there are many alternative ways to carve nature at its joints, there is a mind-independent basis for such carving. A central advantage of this view is that it accommodates the fact that scientific practices may involve different (often incompatible) classifications, facilitating diverse projects and explanatory goals. At the same time, it aims to avoid conventionalism, by providing a metaphysical stance that accounts for the way in which various scientific practices produce knowledge of a mind-independent world (Chakravartty 2007, 2011).

The main challenge faced by the pluralistic realist is to articulate the mind-independent basis for the alternative ways scientists carve up the world. To what sort of objective structures in nature do scientific kinds correspond? Chakravartty appeals to patterns of causal properties in the world, thus providing an account that encompasses both 'essence kinds' and 'homeostatic property cluster' kinds. In what follows, I use his account as a springboard to explicate my own.

I begin by discussing Chakravartty's position that alternative taxonomies are grounded in mind-independent patterns of property distribution. I argue that, while the focus on properties is helpful, the idea that patterns are mind-independent does not accommodate all types of scientific categories. The

difficulty is the requirement that category members share (some or all of) the causal properties that ground a kind. For Chakravartty, this sharing requires a relation of sameness with respect to the relevant individual properties. I argue that in order to accommodate at least some kinds (e.g., biological kinds), the sameness requirement for property sharing should be replaced with a similarity requirement, with respect to individual properties. Scientific kinds are therefore grounded in similarity relations not only in the sense of ‘family resemblance’ – according to which kind members share some, but not all, of the properties that ‘make’ the kinds – but also in a more fundamental sense: sharing a property only requires similarity, not identity. Kinds are grounded in similarity ‘all the way down’.

Next, I draw on insight from psychology to explicate the sort of relation that similarity is. I alleviate the realist’s concerns about similarity, and show that, while similarity is a cognitive construct, similarity-based kinds can still satisfy the criterion for realism. Specifically, knowledge of mind-independent regularities can be extracted from similarity-based scientific taxonomies. I then briefly examine the case study of the concept SYNAPSE, demonstrating that members of the category are similar, rather than identical, with respect to relevant causal properties that characterize the kind, and show that mind-independent regularities among synapses can be extracted from the taxonomy in which the category is embedded. I end by suggesting that the similarity requirement for property sharing better achieves the goals of scientific realism than does the sameness requirement, and discuss the ways in which my approach is distinguished from conventionalism.

Chakravartty, A. (2007): *A Metaphysics for Scientific Realism: Knowing the Unobservable*, Cambridge, UK: Cambridge University Press.

Chakravartty, A. (2011): ‘Scientific Realism and Ontological Relativity’, *The Monist*, 94(2): 157-180.

Representational Practice in Representative Sampling in Public Health

Brandon Boesch – Morningside College, United States

Michaela Schenkelberg – University of Nebraska Medical Center, United States

Philosophical discussions of scientific representation are typically offered primarily to explain how it is that a scientific model is representational (see, e.g., Hughes 1997; Bailer-Jones 2003; French 2003; Giere 2004, 2010; Contessa 2007). While these studies are valuable, there are other contexts in which representational practice is found in science. In this paper, we will analyze a

case study from public health to understand the representational nature of representative samples. Specifically, we will examine the Youth Risk Behavior Surveillance System (YRBSS), used to monitor health-related activities among youth in the United States every other year since 1991 (Brener et al. 2013).

According to the pragmatic view of scientific representation, to understand the representational nature of a vehicle, we must pay attention to the ways in which it is “licensed”—i.e., constructed, developed, and used over time by the scientific community ([removed-for-blind-review]). For the YRBSS, this prominently includes steps taken to develop a sample which matches the demographics of the population and avoids potential biases (e.g., three-stage cluster design and large sample size). But it is also important to understand how the data is meaningful and useful within scientific research. Consider an example from the physical activity sub-category of the YRBSS. In the past thirty years, there has been an enhanced understanding of the dose-response relationship between physical activity and health outcomes, resulting in changes to the YRBSS. The first YRBSS (in 1991) asked students how many days per week they were physically active, but ignored the duration of their activity. Durations were included in 1993, likely to match the (then under-development) recommendations of physical activity guidelines for adolescents (Sallis and Patrick 1994). The items were modified yet again in 2007, when the duration was specified to 60 minutes per day—probably influenced by the (then under-development) 2008 Physical Activity Guidelines (USDHHS 2008). Here we can see that changes in the survey due to theory allow scientists to use the sample to draw inferences about the target population (Suárez 2004).

Another component of the representational nature of representative samples is how scientists constrain their use of the data according to their understanding of the characteristics of the sample. For example, in a 2003 study, scientists examined the relationship between weight status and physical activity, noting that their results may have been skewed by the fact that they used BMI to determine weight status (meaning that some muscular children could be mistakenly classified as overweight) (Levin et al. 2003). In virtue of their awareness of its shortcomings, scientists construct and define the representational uses of the sample—excluding certain targets and limiting confidence in others.

As a tentative account, we will argue that representative samples are representational because of how they are collected and used by scientists, especially in virtue of the (1) methodology of collection, which structures the sampling practice to eliminate risk of bias and sampling errors; (2) the influence of theory in development of a measurement tool; and (3) the constraint of representational aims due to awareness of shortcomings.

- Bailer-Jones, Daniela. 2003. "When Scientific Models Represent." *International Studies in the Philosophy of Science* 17 (1):59–74.
- Brener, Nancy D., Laura Kann, Shari Shanklin, Steve Kinchen, Danice K. Eaton, Joseph Hawkins, and Katherine H. Flint. 2013. "Methodology of the Youth Risk Behavior Surveillance system—2013." *Morbidity and Mortality Weekly Report: Recommendations and Reports* 62 (1):1–20.
- Contessa, Gabriele. 2007. "Scientific Representation, Interpretation, and Surrogative Reasoning." *Philosophy of Science* 74 (1):48–68.
- French, Steven. 2003. "A Model-Theoretic Account of Representation (Or, I Don't Know Much about Art...but I Know It Involves Isomorphism)." *Philosophy of Science* 70 (5):1472–83. <https://doi.org/10.1086/377423>.
- Giere, Ronald N. 2004. "How Models Are Used to Represent Reality." *Philosophy of Science* 71 (5):742–752.
- . 2010. "An Agent-Based Conception of Models and Scientific Representation." *Synthese* 172 (2):269–81.
- Hughes, R. I. G. 1997. "Models and Representation." *Philosophy of Science* 64: S325–36.
- Levin, Sarah, Richard Lowry, David R. Brown, and William H. Dietz. 2003. "Physical Activity and Body Mass Index Among US Adolescents: Youth Risk Behavior Survey, 1999." *Archives of Pediatrics & Adolescent Medicine* 157 (8):816–20. <https://doi.org/10.1001/archpedi.157.8.816>.
- Sallis, James F., and Kevin Patrick. 1994. "Physical Activity Guidelines for Adolescents: Consensus Statement." *Pediatric Exercise Science* 6 (4):302–314.
- Suárez, Mauricio. 2004. "An Inferential Conception of Scientific Representation." *Philosophy of Science* 71 (5):767–79. <https://doi.org/10.1086/421415>.
- United States Department of Health and Human Services (USDHHS). 2008. "2008 Physical Activity Guidelines for Americans." USDHHS. <https://health.gov/paguidelines/pdf/paguide.pdf>

Big data in need of big scientists – Revisiting Polanyi's notion of personal knowledge

Mieke Boon – University of Twente, Netherlands

With the rise of A.I., expert-systems, machine-learning technology and Big Data we may start to wonder whether people as creative, cognitive and intellectual beings will become redundant for the generation of knowledge. Also, the increasing success of machine-learning technology in finding patterns in data makes us ask whether scientific theories will become superfluous. First, it will

be argued that an empiricist conception of science makes these thoughts plausible and provides little room to criticize them. Next, Michael Polanyi's (1958/1962) notion of personal knowledge will be taken as a starting-point in the development of a suitable alternative that avoids the rather defensive position in which an empiricist is forced and preferably does not fall back on naive scientific realism.

Hence, this paper aims at a contribution to critically investigating whether theoretical knowledge and the scientist's role in developing it, will remain crucial – or will arbitrary algorithms, provided by machine-learning technologies for constructing relationships between data-input-output, eventually be able to meet crucial epistemic criteria such as empirical adequacy, reliability and relevance, better than limited humans ever could?

The empiricist strand in the philosophy of science has a long history of making the role of humans in science superfluous, or at least, to downplay their role such as to justify the objectivity of knowledge. Strict empiricism, from Hume to Logical Positivism and anti-realist views such as Van Fraassen, can only assume that, as theory-formation transcends what has been given in empirical data, theories must be understood as just heuristic tools that are only necessary for limited beings. We may also recall the responses to Bogen and Woodward (1988) when claiming that phenomena can be found in data. McAllister (1997) criticizes this view by arguing that there are always infinitely many patterns in any data-set, and so the choice of one as being a phenomenon is subjectively stipulated by the investigator. Glymour's (2000) strict empiricist alternative is that scientists infer from data to patterns by means of statistical analysis, which, according to him, does not involve subjective grounds, and which does not add anything new to the data. Clearly, if, objectively, there is just patterns in data-sets, only subjectively summarized in theories (including 'mini-theories' such as phenomena and laws), it is to be expected that machines will surpass the intellectual contributions of humans.

Polanyi was a physical-chemist, who proposed a philosophical conception of science that accounted for the role of the scientist. Central is his premise that the supposed objective-subjective gap has led to the flawed belief – in positivism and empiricism – that experiences are the objectively reliable part of knowledge, instead of theory. Polanyi defends personal knowledge, in which the contribution of the human intellect to knowledge is crucial. He bridges the objective-subjective gap by claiming that science must be understood as an inherently and inescapably responsible endeavor, because individual scientists necessarily subject themselves to general rules and 'out of passion' take responsibility for convincing others of their theoretical findings. This paper aims to explain the relevance of Polanyi's argument for better understanding the indispensable role of the human intellect.

Gender and the Measurement of Fertility: A Case Study in Critical Metrology

Marion Boulicault – MIT, United States

Human fertility is in an apparent state of crisis. In July 2017, scientists reported that sperm counts among men from North America, Europe and Australia have decreased by 50 – 60% since 1973, with no sign of halting (Levine et al. 2017). For women, the story is bleak and familiar: women’s fertility decreases with age, yet women are waiting longer and longer to have children (Kincaid 2015). Undergirding these crisis narratives is an unstated assumption: fertility is measurable. That is, scientific reports that fertility is declining presuppose that it’s possible to successfully measure and compare fertility diachronically. In this paper, I investigate this assumption by examining the practice of fertility measurement, i.e. the standards, methods and instruments by which the phenomenon of fertility is quantified. By comparing two current gold standard fertility measures – semen analysis in men, and ovarian reserve testing (ORT) in women – I argue that cultural ideas about gender play a significant role in constructing fertility as a measurable phenomenon. Different temporal assumptions implicit in semen analysis and ORT reflect and enforce different gendered imperatives of responsibility over reproduction. More specifically, temporal dimensions implicit in the measurement of fertility in men and women supports a view of women as more responsible for – and therefore more to blame for – infertility. I conclude by arguing that, in the case of semen analysis and ovarian reserve testing, it’s not just fertility that’s being measured, but also degrees of adherence to traditional Western norms of masculinity and femininity, and, more precisely, to cultural responsibilities over motherhood and fatherhood.

This paper also has a methodological aim. Important philosophical work has been done in investigating measurement as a metaphysical and epistemological phenomenon (see Tal (2017) for an overview). This paper is in part a call to attend more closely to the political and ethical dimensions of metrological practices. What role does measurement play in the creation and maintenance of social norms and, conversely, how are social norms reflected in our measurement practices? How can we best conceptualize and investigate the intersection of measurement and oppression? Answering these questions, I contend, requires an interdisciplinary approach that pays close attention to measurement as a material and social practice. This approach – which I propose we call ‘critical metrology’ – is what I attempt to demonstrate in my analysis of fertility measurement.

Campoamor, Danielle. 2017. "Thanks, But Meghan Markle's 36-Year-Old Uterus Does Not Need Your 'Concern.'" *Romper*. November 27, 2017. <https://www.romper.com/p/thanks-but-meghan-markles-36-year-old-uterus-does-not-need-your-concern-5555914>.

Kincaid, Ellie. 2015. "American Women Are Waiting Longer than Ever to Become Mothers." *Business Insider*. June 15, 2015. <http://www.businessinsider.com/average-age-of-mother-having-first-child-going-up-2015-6>.

Levine, Hagai, Niels Jørgensen, Anderson Martino-Andrade, Jaime Mendiola, Dan Weksler-Derri, Irina Mindlis, Rachel Pinotti, and Shanna H Swan. 2017. "Temporal Trends in Sperm Count: A Systematic Review and Meta-Regression Analysis." *Human Reproduction Update* 23 (6):646–59. <https://doi.org/10.1093/humupd/dmx022>.

Tal, Eran. 2017. "Measurement in Science." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2017. Metaphysics Research Lab, Stanford University.

<https://plato.stanford.edu/archives/fall2017/entries/measurement-science/>.

Teaching Conceptual Change: Can Building Models Explain Conceptual Change in Science?

Dragana Bozin – University of Oslo, Norway

This paper considers how novel scientific concepts (concepts which undergo a radical conceptual change) relate to their models. I present and discuss two issues raised by respectively Chin and Samarapungavan (2007) and Nersessian (1989) about perceived (and persistent) difficulties in explaining conceptual change to students. In both cases models are either seen as secondary to concepts/conceptual change or seen as inessential for explanation. Next, I provide an example which to some extent counters these views. On the basis of that example I suggest an alternative view of the role of models in conceptual change and show that the latter could have beneficial implications for teaching conceptual change. The example in question is Robert Geroch's modeling of Minkowski spacetime in *Relativity from A to B* (1981).

It seems reasonable to think that understanding the conceptual transformation from space and time to spacetime first, makes it easier to build a model of spacetime. This is the underlying assumption that Chin and Samarapungavan make (2007). Their objective is to find ways to facilitate conceptual change because they see the lack of understanding of the conceptual change that produced the concept as the main obstacle for students' ability to build a model of it. I argue that this is not necessarily the case: in certain cases

(spacetime for example) building the model can facilitate understanding of the conceptual change.

In a similar vein, although understanding how scientific concepts developed can often give clues for how to teach them I argue that in some cases the historical approach is counterproductive. Nersessian argues that the same kind of reasoning used in scientific discovery could be employed in science education (Nersessian, 1989). I essentially agree with this view but with a caveat. I argue that in some cases the historical approach might be constraining and in particular that the spacetime example shows that ignoring the historical path in certain cases is more successful.

Additionally Geroch's way to model spacetime can be of consequence for teaching relativity and quantum mechanics to high school students. Physics is traditionally taught through solving equations and performing experiments which is ill suited for relativity and quantum mechanics. Norwegian curriculum requirements include that students be able to give qualitative explanations as well as discuss philosophical and epistemological aspects of physics. According to ReleQuant (University of Oslo and the NTNU project on developing alternative learning resources for teaching relativity and quantum mechanics to high school students) this opens the door for introducing qualitative methods in teaching high school physics. The conclusion that ReleQuant draws from this is that historical approaches may be profitable when teaching quantum physics on the high school level.

The historical approach might not always be effective – as it is not in teaching spacetime. Teaching through building a model “from scratch” might work better. Building a model from with no or little reference to theory could be viewed as a qualitative method and would essentially be in agreement with the overall ambition of the ReleQuant project.

How to Philosophically Tackle Kinds without Talking About ‘Natural Kinds’

Ingo Brigandt – University of Alberta, Canada

Traditional visions of natural kinds as proposed in analytic metaphysics, which may assume that a kind is characterized by an intrinsic essence, have been found to be wanting, because they fail conform to the diversity of kinds found in biology and other special sciences. In contrast, philosophers of science have put forward alternative accounts, which endeavor to capture natural kinds as they are actually discovered in science. Examples are Boyd's notion of kinds as homeostatic property clusters, Khalidi's kinds as nodes in causal networks, Slater's stable property clusters, and Franklin-Hall's kinds as categorical

bottlenecks. However, each of these philosophical accounts of natural kinds is still too restrictive, because while covering many cases the account fails to capture other instances of bona fide kinds by—as I argue—failing to capture what makes these kinds scientifically important. Therefore, it seems misguided to endorse or even attempt to put forward a unique, general philosophical theory of what qualifies something as a natural kind (as opposed to a nominal kind). Instead, philosophers might better investigate kinds on a case-by-case basis, with close attention to how such a particular kind figures in scientific theorizing and practice.

While being sympathetic to this latter deflationary attitude toward philosophical theories of natural kinds, in this talk I will put forward some general points on how to philosophically discuss and investigate kinds after all. My core tenet is that any grouping of objects into an alleged kind can only be assessed with reference to the interests and aims of those employing the kind (and based on whether these aims are legitimate). In contrast to several other accounts (e.g., Khalidi's), I eschew prioritizing epistemic over non-epistemic aims, and instead argue that genuine kinds may answer to a combination of epistemic, practical, moral, and political interests. This argument could be made in the cases of race and sex, which are kinds that not only exhibit a causal interplay of biological and social features, but where any defensible account of what race or sex is must also answer to political concerns. In this talk, though, my argument will focus on psychiatric kinds, in particular personality disorders. I will argue that some psychiatric kinds do not just answer to practical considerations related to medical treatment, but also to other normative and social concerns.

This vision of kinds as essentially [sic!] answering to human interests, including various non-epistemic aims, is a far cry from the agenda of analytic metaphysics to discuss natural kinds in order to articulate a mind-independent structure of reality or to exhibit the fundamental level of reality. But this is exactly why philosophers of science cannot cede the field and have to continue with general discussions of 'kinds'. At the same time, I argue that the term 'natural kind' should be avoided, because it erroneously suggest that all kinds are physico-chemical or narrowly biological (as opposed to involving contingent social features) and that kinds are merely part of some order of nature (as opposed to answering to human interests).

Public Epistemic Trustworthiness & Lay Participation: The Case of Psychiatric Classification

Anke Bueter – Leibniz Universität Hannover, Germany

Taxonomic decision-making in psychiatry is highly controversial, as could be witnessed again during the latest revision of the DSM (APA 2013). For one, this controversial nature stems from the state of knowledge incorporated in such nosologies that many find insufficient. For another, it is due to the enormous practical consequences of nosological changes in the DSM (or the ICD, which is used in most European countries): they impact the course of research as well as reimbursement policies, educational practices, and the treatment of patients. The combination of these points has led to an abundance of DSM-critiques and a severe lack of credibility of psychiatric classification in the public eye.

My paper deals with the question of how the public epistemic trustworthiness of the DSM might be enhanced. In particular, I argue that an increased integration of patients, advocates, and care-givers – which has been rejected as “politically correct nonsense” before (Spitzer 2004) – is helpful in this respect. The latest DSM revision has made some first steps in this direction, which can be justified as well as enhanced.

According to Wilholt (2013), public epistemic trustworthiness requires a distribution of inductive risks that figures in the public’s expectations. In the case of psychiatric classification, it is plausible to assume that the public’s expectation is to minimize risks for patients, rather than for other stakeholders such as the pharmaceutical industry or health insurers. I argue that this can be generalized as a requirement to make value-laden decisions in a manner that represents patient’s best interests. Such a representation is a necessary, yet insufficient condition for public epistemic trustworthiness: In addition, it requires a reason for the public to believe that patients’ values and interests are in fact appropriately represented (Irzik & Kurtulmus forthc.).

I will then show that psychiatric classification involves value-judgments at several points. For example, this can concern decisions on the disorder-status of conditions or behaviors and the weighing of associated risks. As taxonomic decisions always trade between risks of over- versus underdiagnosis, the perspective of patients is a relevant input regarding whether it would be better to err on the side of being too rigid or too inclusive in the criteria for particular mental disorders.

Based on this, I will argue that in the case of psychiatric classification, patients’ values are best to be represented by patients and advocates themselves rather than by scientific experts. This is due to the special social situation of

psychiatric classification as facing a long history of public distrust. Not only is this distrust expressed by many organizations and movements such as critical psychiatry, psychiatry survivors, or neurodiversity, it also has a long (academic and popular) tradition with roots in the anti-psychiatry movement of the 1970s. Thinking about the historic track-record of psychiatry and psychiatric care (e.g. regarding the treatment of homosexuals or hysteric women, forced hospitalizations, lobotomies, etc.), this distrust is also not completely unreasonable. In a nutshell, the representation of patients' interests and values by scientists is exactly what is questioned by the public – and therefore fails to provide an appropriate reason to ground public epistemic trustworthiness.

Finally, I will briefly discuss different models of patient integration in hindsight of their advantages and potential pitfalls, for example regarding potential conflicts of interests or the representation of dissent among advocates.

American Psychiatric Association (APA) (2013). *Diagnostic and Statistical Manual of Mental Disorders, Fifth Edition: DSM-5*. Washington, DC: APA.

Irzik, G. & Kurtulmus, F. (forthc.). What is Epistemic Public Trust in Science? *British Journal for the Philosophy of Science*.

Spitzer, R. L. (2004). "Good Idea or Politically Correct Nonsense?". *Psychiatric Services* 55: 1.

Wilholt, T. (2013). Epistemic Trust in Science. *British Journal for the Philosophy of Science* 64: 233-253.

Pluralism and Representation in Biology

Daniel Burnston – Tulane University, United States

One of the most-pursued debates in current philosophy of biology and neuroscience is whether mechanistic explanations exhaust explanations—i.e., whether there are system properties that are explained by fundamentally non-mechanistic frameworks. Some candidate frameworks that are touted as non-mechanistic are dynamical state space models (Ross, 2015), and network models (Huneman, 2010). Proponents of limited mechanistic explanation think these models are independently explanatory. Mechanists counter by arguing that the alternative models are abstractions of, or must map onto, more standard mechanistic explanations (Kaplan & Craver, 2011; Levy & Bechtel, 2013).

From the framework of integrative pluralism about explanation (Mitchell, 2003), these debates are odd. Pluralists argue for the necessity of multiple types of representations in explaining the behavior of complex systems, but

also cite the need for these different models to relate to each other in the course of explanations. Integrative pluralism requires an account of what makes representations integrate successfully, and this account is not yet fully developed. A notion of integration should explain exactly how distinct representations contribute to an explanation, and must make sense of (i) the necessity of multiple representations, (ii) why representations need to be integrated, and (iii) how they do so.

I argue that the best account of integration can be gained by taking an inferentialist approach to representation (Suárez, 2004). Unlike similarity- or isomorphism-based views of scientific explanation, inferentialist views argue representations contribute to explanation by affording and constraining the inferences and hypotheses that practicing scientists can make about the system. I articulate a version of inferentialism on which a given representation entails certain inferences and forbids certain inferences, but also remains non-committal about a range of other possible inferences. I discuss examples on which novel kinds of representation, for instance phase-space models, are introduced to make explicit inferences that are not entailed by mechanistic or data models.

I then introduce the concept of inferential closeness. A representation A is inferentially closer to a system property P than a representation B if A explicitly entails an inference about P that is consistent with but not entailed by B. This makes sense of why we have better overall explanations when multiple representations are integrated, but also why scientists need to fill inferential gaps with new representations. I conclude by discussing the ramifications of inferentialism for types of pluralism and for debates about the sufficiency of mechanism.

Huneman, P. (2010). Topological explanations and robustness in biological sciences. *Synthese*, 177(2), 213-245.

Kaplan, D. M., & Craver, C. F. (2011). The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective. *Philosophy of Science*, 78(4), 601-627.

Levy, A., & Bechtel, W. (2013). Abstraction and the Organization of Mechanisms. *Philosophy of Science*, 80(2), 241-261.

Mitchell, S. D. (2003). *Biological complexity and integrative pluralism*: Cambridge University Press.

Ross, L. N. (2015). Dynamical models and explanation in neuroscience. *Philosophy of Science*, 82(1), 32-54.

Suárez, M. (2004). An inferential conception of scientific representation. *Philosophy of Science*, 71(5), 767-779.

Against the Hierarchical View of Theories

Julia Bursten – University of Kentucky, United States

To butcher a line from Nancy Cartwright: There is view about inter-theory relations that is so deeply entrenched that it doesn't even have a name of its own. It is the view that scientific theories are ordered along a spectrum of generality or fundamentality, ranging from the most fundamental and most general laws of physics to the least fundamental or most particular, context-sensitive laws of various social sciences. Less fundamental theories are often said to be more complex than more fundamental theories, and more fundamental theories are said to be more wide-ranging. Physics, the discipline responsible for the most general and most fundamental theory, applies to all biological systems, whereas biology, a more complex and less fundamental theory, does not apply to all physical systems, to say nothing of psychology or economics. The more general a theory is, the more fundamental it is. Likewise, within a discipline, theories are ordered by fundamentality: in physics, quantum mechanics is more fundamental than continuum mechanics, which is more fundamental than theories of turbulent fluid flow. Fundamentality may be, but is not necessarily, a theoretical virtue; whether or not fundamentality breeds better theories, or just more general ones, is, in a sense, behind much of the dispute between reductionists and emergentists.

I call this view the hierarchical view of theories. While the hierarchical view may be more apparent on the surface of reductionist approaches to inter-theory relations, it underscores most emergentist accounts as well; and where it is absent, the exception proves the rule, as the resulting accounts appear rather unlike traditional emergentist views. My main aim in this talk is to show that the hierarchical view of theories has led philosophers of science on both sides of the reduction–emergence debate to an impoverished infrastructure for understanding inter-theory relations. By focusing too narrowly on the explanatory and predictive work accomplished at individual or component levels, the hierarchical view excludes the epistemic contributions of the conceptual strategies employed to connect higher-level theories to lower-level ones. These strategies are an essential and as-yet ill-understood piece of architecture in the epistemology of science, and the hierarchical view has occluded them from analysis. I demonstrate this occlusion through a critical re-examination of an example, introduced by Eric Winsberg, of multiscale modeling of nanoscale crack propagation. My analysis attends to the particularities of how models and theories connect across different levels or frameworks. I show that this account leads to a more fruitful and robust picture of inter-theory relations that can explain how models and theories

support epistemic and explanatory activities in science.

My account emphasizes the roles of algorithms, heuristics, mathematical and computational processes, and other conceptual strategies in stitching together various theoretical frameworks. The result is a view of inter-theory relations that steers philosophical consideration toward the importance of the work accomplished at the borderlands between one theoretical framework and its neighbors.

Individuating Genes as Types or Individuals

Ruey-Lin Chen – National Chung Cheng University, Taiwan

This paper explores individuation of genes from the perspective of scientific practice. It thus involves both the issue of biological individuality and the issue of approaches from theoretic constructions or experimental practices. I argue that the transgenic technique can individuate “a gene” as an individual while the technique of gene mapping in classical genetics can only individuate “a gene” as a type or a kind. Herein we find double extensions of the term “a gene”. Thus, I also discuss this semantic phenomenon in using “a gene”.

To date, “what is a gene?” and other related questions have been asked once again by philosophers, historians, and scientists of biology (Kitcher 1992; Falk 2010; Carlson 1991; Maienchein 1992; Portin 1993; Waters 1994, 2007; Beurton, Falk, and Rheinberger 2000; Snyder and Gerstein 2003; Stotz and Griffiths 2004; Pearson 2006; Reydon 2009; Baetu 2012). Those questions were frequently embedded in the discussion about both the definition of “gene” and the gene concept. Philosophers have disputed and continued to dispute on whether or not there is a single or united definition or concept of gene. In spite of few exceptions (e.g., Waters), however, most of them approach to this issue from the theoretic perspective.

This paper explores experimental individuation of genes along an alternative direction (i.e., the transgenic technology), exploring the possibility that a gene is individuated as an individual. The question of what a gene is explicitly presupposes the problem of the gene individuality; and identifying a gene presupposes individuating the gene. According to the literature of analytic metaphysics, “individuation” is traditionally understood in a metaphysical and an epistemological sense. Beuno, Chen, and Fagan (2018) add a practical sense to the term and connecting the three senses interpret to “individuation”. They characterize “individuation” and “individuals” as “an individual emerges from a process of individuation in the metaphysical sense. Epistemic and practical individuation, then, are processes that aim to uncover stages of that metaphysical process.” (Beuno, Chen, and Fagan 2018, in production) The

approach to individuation of genes adopted in this paper follows their characterization, especially focusing on the process of epistemic and practical individuation.

This paper argues for the following two points: (1) Classical geneticists have individuated a gene as a type. (2) The experiments using the transgenic technique can individuate a gene as a particular or an individual. The two points indicate two different kinds of individuation: individuation of a type and individuation of an individual. One may wonder whether or not “the individuation of a type” is an inconsistent phrase. In order to answer this question, I will discuss in what sense we individuate a type. The discussion thus involves the relationship between kind and individual in the context of experimentation.

Baetu, Tudor M., 2012. “Genes after the Human Genome Project.” *Studies in History and Philosophy of Biological and Biomedical Science*, 43: 191-201.

Beurton, P., R. Falk, and H.- J. Rheinberger, 2000. *The Concept of the Gene in Development and Evolution: Historical and Epistemological Perspectives*. Cambridge, UK: Cambridge University Press.

Beuno, Otavio, Ruey-Lin Chen, and Melinda B. Fagan, 2018. “Individuation, Process, and Scientific Practice.” In Otavio Beuno, Ruey-Lin Chen and Melinda B. Fagan (eds). *Individuation, Process, and Scientific Practice*. New York: oxford University Press. (In production)

Falk, Raphael, 2010. “What is a gene – revised” *Studies in History and Philosophy of Biological and Biomedical Science*, 41: 396-406.

Kitcher, P. S., 1992. “Gene: current usages.” In E. Keller and L Lloyd (eds.), *Keywords in Evolutionary Biology*, Cambridge, MA: Harvard University Press, pp. 128-131.

Pearson, Helen, 2006. “What is a Gene?” *Nature*, 441(25): 399-401.

Portin, P. (1993). The concept of the gene: Short history and present status. *Quarterly Review of Biology*, 68(2): 173-223.

Reydon, Thomas, 2009. “Gene Names as Proper Names of Individuals: An Assessment.” *British Journal for the Philosophy of Science*, 60(2): 409-432.

Snyder, Michael and Mark Gerstein, 2003. “Defining Genes in the Genomics Era.” *Science*, 300(5617): 258-260.

Stotz, Karola and Paul Griffiths, 2004. “Genes: Philosophical Analyses Put to the Test.” *History and Philosophy of the Life Sciences*, 26: 5-28.

Waters, Kenneth C., 1994. “Genes made molecular,” *Philosophy of Science*, 61: 163–185.

Waters, Kenneth C., 2007. “Molecular genetics,” *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/molecular-genetics/>

When should we accept that a phenomenon doesn't exist? Memory transfer, conflicting evidence, and defeasibility

David Colaço – University of Pittsburgh, United States

The identification of phenomena is a critical scientific research activity, as it is responsible for the discovery and characterization of the types of events to be explained by theory. To fulfill their theoretical and practical aims, researchers set out to accept characterizations of phenomena when empirical findings are put forward in their favor. When a characterization of a phenomenon is accepted, researchers theorize and experiment in a way that is consistent with the existence of the phenomenon. However, many episodes in the history of science involve the abandonment of characterizations of phenomena that were once empirically promising. This raises a question: under what circumstances do researchers reject the existence of a characterized scientific phenomenon, despite evidence that appears to support it?

I explore a case in which the existence of a phenomenon was rejected: the research on memory transfer. This alleged phenomenon was described as the transfer of learned behavior by the insertion of tissue from a trained donor organism to an untrained receiver. It received a great deal of attention from scientists and the public alike, due to its implications and to researchers' use of sensational experiments involving cannibalism. Formulated and defended in light of empirical findings, the characterization of memory transfer was considered by some to be accurate. This led to a cottage industry about its characterization, its theoretical significance, and its underlying mechanistic details. The research program was abandoned, and contemporary scientists generally consider the "phenomenon" to not exist.

Historians of science have questioned the motives of the scientific community that abandoned research on memory transfer. For instance, Harry Collins and Trevor Pinch argue that there was no "decisive technical evidence" that disproved the existence of memory transfer, and that research was abandoned due to disinterest in the purported phenomenon. They base their argument on the fact that no challenge to memory transfer applies to all experiments whose findings were thought to provide support for the characterization of the alleged phenomenon.

I challenge Collins and Pinch's argument by illustrating why researchers were justified in abandoning memory transfer, which is based on the defeat of the evidence provided in the characterization's favor. By defeating all evidence, any empirically motivated reason for accepting the existence of a phenomenon as characterized is eliminated. With this strategy, researchers do not simply provide evidence to challenge a characterization of the phenomenon; they also

demonstrate the faultiness of the experiments whose findings are thought to support the characterization in the first place. My response answers the question of why there was good reason to reject memory transfer. My analysis also shows how, generally, researchers can provide reason to reject a scientific claim when there appears to be conflicting evidence, even if they do not have an alternative theoretical explanation for some of the evidence that has been collected.

Explainable AI and Scientific Explanation

Kathleen Creel – University of Pittsburgh, United States

Machine learning systems have become increasingly sophisticated and powerful within the last ten years. Software engineers and scientific researchers have capitalized on advances in algorithms and computing power to classify particle collisions at the LHC, detect cancer with artificial neural networks, or identify fossil pollen with visual recognition software (Tcheng et al. 2016; Esteva et al. 2017). But despite these advances in prediction and classification, complex machine learning systems often do not provide explanations for their decisions. Their opacity has prompted a new area of research: explainable AI.

Practitioners in explainable AI have identified three interrelated goals: explanation, interpretability, and justification. Interpretability is often defined as human understandability of the functioning of the program, whether produced by inspection of code, analysis of the algorithm, or output generated by the program. Justification provides an explanation of “why a decision is a good one, but it may or may not do so by explaining exactly how [the decision] was made” (Biran and Cotton 2017). Explanation is used in many ways, and its definition seems to depend on the purposes of the researchers and their users. In applied machine learning, success is measurable by ease and frequency of use among the constituent population. For example, when doctors do not trust decision-assisting software, they will not use it in clinical practice, even when they are told that its use decreases diagnostic errors (Hutson 2017). Since studies showed that doctors value explanations for the diagnosis very highly, especially step-by-step rationales for decisions, the developers of these systems sought to provide explanations of the kind desired (Ye and Johnson 1995; Symeonidis et al. 2009).

However, when machine learning is used in the practice of science, there is more to be said about what epistemic goods are required. This is especially true for unsupervised machine learning algorithms, which are often fast, powerful, and opaque. While supervised machine learning can sometimes be

taught to recognize objects or images using criteria similar to the ones that might be used to train a human lab technician, unsupervised learning relies on the algorithm to develop its own categories and variables for classification. This leaves the scientist without an explanation for why certain classifications were made rather than others, or what makes the groupings discovered similar. This can be especially problematic when the algorithm is performing a task previously done with human perception or cognition.

In this paper, I bring resources from the philosophical literature on scientific explanation into conversation with the questions raised by explainable AI. When machine learning is used in scientific contexts for classification or prediction, what counts as a good explanation of its functioning? Is justification enough, or ought scientists to require that algorithms be transparent or interpretable? (Ananny and Crawford 2016; Burrell 2016) Using the interventionist framework, I present a principled distinction between scientific explanations that require that programs be interpretable and those for which justification is sufficient. This adds precision to the debate over the meaning, proper functioning, and political importance of algorithmic transparency.

Epistemic Responsibility in Science

Haixin Dang – University of Pittsburgh, United States

Responsibility is a central concept in the social epistemic practices of science, but the concept has often been left unanalyzed. Modern science is conducted by teams, sometimes numbering in the hundreds or even thousands. The latest physics article on the mass of the Higgs boson had over 5,000 listed authors. To what extent are these authors epistemically responsible for the discovery of the mass of the Higgs boson? We need a concept of epistemic responsibility which can ground our determination of who should get acknowledgment or rewards for a scientific discovery and also who should be sanctioned when a scientific claim turns out to be false or faulty.

In face of collaboration, one may be tempted to deny that epistemic responsibility is truly possible or tenable. Huebner, Kukla, and Winsberg (2017, 2014) have indeed argued that an agent is epistemically responsible if and only if she is able to give a consistent “justificatory story” for her assertions. A justificatory story involves reasons for epistemic standards employed, methodological choices taken, and methodical judgments made. I call this a unified view in contrast to my tripartite account because HKW do not take epistemic responsibility to be further conceptually decomposable. I argue that an unified account is untenable for collaborative science because disagreement is inherent in large social groups, which makes the requirement for one

consistent “justificatory story” to undergird every scientific claim extremely difficult.

I develop instead a tripartite account of epistemic responsibility which can help us locate responsibility in collaborative science. I argue that epistemic responsibility in science has three distinct senses: attributability, answerability, and accountability. The account can be briefly summarized as follow: An agent can be epistemically responsible for a claim that P if that claim can be properly attributed to the agent. An agent can be answerable-responsible for a claim that P in so far as the agent is able to report the reasons and justifications for holding that P. And finally an agent can be held accountable for a claim that P if it is appropriate to blame or praise the agent for asserting that P in accordance to epistemic norms. This account of epistemic responsibility is analogous to an existing account of moral responsibility which have been advanced by David Shoemaker.

Finally, I develop further one aspect of my account: answerability. Collaborators can be answerable for different parts of an epistemic project, however this also means that collaborators may give different sets of reasons and justifications for holding that P that do not match up with each other. In response to this worry, I argue that some degree of disagreement among collaborators over the reasons must be tolerated in collaborative science so long as there is consensus over the conclusion or final discovery claim. This may be a surprising result, but I argue that an account of collective justification in science must allow for certain kinds of disagreements. I reject the HKW view which require agreement for both reasons and conclusions.

Analyzing the meaning of causal claims in genetic epidemiology: a pragmatic point of view

Leen De Vreese – Ghent University, Belgium

What is the role of genes in the causation of disease? In the second half of the previous century, the discovery that variations in singular genes were responsible for the development of some so-called “monogenic diseases” led to the hope of finding much more disease-causing genes in the years to follow. However, things turned out not to be so simple. Instead of finding straightforward genotype-phenotype linkages, further research foremost uncovered the complexity of “genetic diseases”, even of those that were formerly supposed to be of the simple, monogenic kind. On the other hand, the progress of science resulted in the finding that influencing genetic factors could be detected for diseases of all kinds, not just for “genetic diseases” as such. Hence, the awareness grew that – in the end – any disease is somehow

influenced by genetic factors, and therefore can be claimed to be genetic in some sense. This led to a paradoxical situation, in which the criteria for the distinction between genetic and nongenetic diseases faded away.

In the philosophical literature, attempts have been made to resolve this paradoxical situation and to pinpoint ways to make a clear conceptual distinction between genetic and nongenetic diseases. Most of the proposed solutions focus on the very general question of what justifies the claim that a gene is the main cause of a certain disease. This means that the problem is framed in terms of a causal selection problem.

However, the solutions have been criticized and have been turned away from in the recent literature, in which the focus lies on network and interaction approaches, situating the role of genes in the whole, complex causal network.

Although these recent approaches seem more elegant to sketch an overall view of the role of genes in the whole of the development of a disease, and are therefore preferable to those solutions that try to make a strict distinction between genetic and nongenetic diseases, they still do not give us much insight in how to interpret specific scientific claims about genetic causes of disease.

Starting from a focus on practice rather than theory, I will therefore argue for a switch in the debate by focusing on the analysis of the meaning of causal claims citing (different kinds of) genetic causes of disease, rather than on pursuing an overall description of what a “genetic disease” is. In doing this, I also join in the pragmatic view of Lisa Gannett (1999) who states that genetic explanations are always dependent on the context of the explanation and the epistemic interests that are involved, and hence never totally objective in the sense of being devoid of pragmatic content.

I will build further on my former work, in which a conceptual framework was developed which consists of a system of causal concepts with different “strengths”. In my talk, I will analyse the meaning of different kinds of causal claims from genetic epidemiology using this conceptual framework. In doing this, I will elaborate a much more concrete version of the pragmatic approach. Also, by using the framework to structure our knowledge about the genetic causes of disease I will be able to make a start in clarifying the conceptual mess in scientific talk about genetic disease causation.

Gannett Lisa (1999). ‘What’s in a Cause?: The Pragmatic Dimension of Genetic Explanations,’ *Biology and Philosophy* 14: 349-374.

The Practical Use of Intentional Information Concepts in Animal Behavior Research

Kelle Dhein – Arizona State University, United States

Philosophers working on the problem of intentionality in non-linguistic contexts often invoke the concept of biological fitness as the objective grounds for attributing intentionality to living systems. Ruth Millikan (1984), for example, has argued that intentionality supervenes on evolutionary history such that the intentional content of a sign is the product of the biological functions that sign has mediated in the past. More recently, Brian Skyrms (2010) has characterized intentionality using evolutionary game theory, arguing that intentional relationships naturally arise within sender-receiver systems of agents that have fitness values, observe environmental states, and communicate those observations with each other. However, these biological theories of intentionality don't square with the way experimental biologists in the field of animal behavior research use intentional concepts to guide their laboratory research. In this paper, I take the stance that philosophers can make progress on longstanding philosophical puzzles, like that of naturalizing intentionality, by carefully analyzing successful scientific practice. To demonstrate that claim, I analyze the way biologists use intentional information concepts in the highly successful program of eusocial insect navigation research.

I argue that (1) animal behavior researchers hang intentional information concepts on goal-directed function, not the deep history of natural selection, and (2) that the intentional information concepts deployed by animal behavior researchers have practical value in experimental contexts. More specifically, I argue that animal behavior researchers use intentional information concepts as placeholder relations to be filled in with causal details through empirical work. A properly functioning ant in a natural environment constitutes a complex web of dynamic causal interactions. In principle, researchers could answer any question about how ants achieve goal Z by identifying and articulating specific causal chains of that web. Because of the way animal behavior researchers hang intentional relationships on goal-directed function, the behavioral goal under investigation constrains what causal configurations can potentially fill in that relationship. Whatever the causal details are for how ants realize Z, those causal details need to mesh with other facts about how ants realize Z. If, for instance, Z is successfully navigating home after a winding foraging run, the causal story of how ants achieve Z needs to account for how ants are able to achieve Z in such a wide variety of circumstances and environments.

In debates over the utility of biological information concepts, Sahotra Sarkar (1996, 2000) has argued that the concept of information failed to gain a substantive role in 1960's molecular genetics because informational approaches to genetics failed and informational theories about genetics turned out to be false. In light of my argument for (2), I conclude by arguing that unlike in molecular genetics, intentional information concepts play a substantive role in animal behavior research.

Millikan, R. G. (1984). *Language, thought, and other biological categories: New foundations for realism*. MIT press.

Sarkar, S. (1996). *Biological Information: A Skeptical Look at Some Central Dogmas of Molecular Biology*, in S. Sarkar (ed.), *The Philosophy and History of Molecular Biology: New Perspectives*. Dordrecht: Kluwer, 187-231.

Sarkar, S. (2000). *Information in genetics and developmental biology: Comments on Maynard Smith*. *Philosophy of Science* 67: 208-13.

Skyrms, B. (2010). *Signals: Evolution, learning, and information*. Oxford University Press.

Representing Biology as Process

John Dupré – University of Exeter, United Kingdom

Gemma Anderson – University of Exeter, United Kingdom

James Wakefield – University of Exeter, United Kingdom

This session developed out of an interdisciplinary project involving an artist (Gemma Anderson), a cell biologist (James Wakefield) and a philosopher (John Dupré). The overall goal of the project is to develop better ways of representing biological processes. It also continues earlier work by Anderson (2017) on drawing as a way of knowing. As Wakefield will argue, the decline of drawing as a practice in biological research has had deleterious consequences for some aspects of biological research.

In the first stages of the project, Anderson and Wakefield have worked together to produce images of mitosis, the process that is the central research topic of the Wakefield lab. The aim was to produce a two-dimensional image that somehow represented the full sequence of transitions involved in mitosis. A method was developed that translated Wakefield's understanding of mitosis into an image in which the vertical dimension represented time, and a number of features (colour, thickness of line, distance from the centre, etc.) represented crucial aspects of the mitotic process. A number of different images have been generated, representing different organism with differences

in their processes of mitosis. The process of producing these images and its rationale will be presented in fuller detail during the session.

The outcome of this activity has been a series of images, which we refer to—for reasons that will be obvious on seeing them—as mitosis pots. These are, perhaps needless to say, very different from familiar textbook depictions of mitosis. The causal connection between features of mitosis and features of the mitosis pots gives us confidence that the images could be used to represent common features and specific differences in mitosis, though it is also clear that such use would require a degree of training. We shall discuss the costs and benefits of introducing such images into the practice and teaching of biology. Preliminary exploration of this question has involved soliciting reactions from other scientists.

We felt that a standard format of three 30 minute presentations representing distinct disciplinary perspectives would not do justice to the interdisciplinary nature of the project. Instead, we propose three brief introductions to aspects of the project by each of the participants, followed by a more detailed presentations of the first completed subproject, which is collaborative work by Anderson and Wakefield. We will allow some time for discussion at the conclusion of each part of the session. Dupr will also chair a concluding 15 minute discussion session.

Abstracts of Presentations

1. “From Process to Representation”, (Dupré, 15 minutes).

In this general introduction I shall describe very briefly the process ontology of biology that I have been developing over the last ten years or so. This project has also led me to think about ways in which a process ontology both seems obvious to many philosophers and yet at the same time unnecessary. One such reason is that the standard representation of biological phenomena in cartoons of boxes and arrows looks already processual enough: the arrows represent change, after all. However, I think that these images can be very misleading in representing biological process. First, and less importantly, the images are flat, in the sense that everything appears to be happening at the same time. Ideally, a sequence of arrows in, say, a metabolic process, should emerge from the paper along a time dimension at right angles to the paper.

Perhaps that is obvious. Much more important, representations of this standard kind may confirm the assumption that the relata they depict are stable, autonomous objects. But in fact, or so I argue, they are themselves processes, dynamically stabilised by a range of processes, perhaps even including that being represented. Such reflections encouraged my discussion with Gemma Anderson of more dynamic ways of representing the “objects” themselves, which led to the present project.

2. “The dynamic process of mitosis: changing experimental practice in relation to the observation-participation axis”, (Wakefield, 15 minutes).

Mitosis is a fundamental biological process in which living cells re-organise protein filaments, termed microtubules, into a complex, dynamic structure - the mitotic spindle - within which chromosomes can be moved and segregated, ultimately facilitating cell division. The archetypal structure of the mitotic spindle was first described to the scientific community by Walther Flemming in a series of elegant, detailed and descriptive, but static, drawings in 1882. Recent advances in live imaging have allowed the “dissection” of distinct molecular pathways that contribute to spindle formation through the spatio-temporal re-organisation of microtubules, greatly assisting our understanding of this process.

However, the decline of drawing in scientific practice is not without consequence. Whereas the training of cell biologists was previously centred around direct participation, using drawing to synthesise many thousands of direct microscope-based observations of live and “fixed” cells, technological and methodological advances in live cell imaging, image acquisition and quantitative image analysis now shift the scientist to the position of observer. Though the validity of this transition can be justified in terms of bias removal and detailed measurement of molecular pathways, we argue that the decline of exploratory imagination as part of the scientific process limits the richness of hypothesis formation.

3. “From drawing as epistemology to representing biology as process”, (Anderson, 15 minutes).

Anderson’s artistic and academic work has investigated the epistemic costs of the decline of graphic skills in the Life Sciences (Anderson, 2014). Anderson will introduce her explorations into the epistemic value of drawing in the context of biology and its role in articulating a 'representational grammar' (Kress and Leewen, 2006) to facilitate reflection on, and understanding of complex scientific concepts that linguistic description alone can struggle to support in the same way (Anderson 2014; Anderson 2017). Anderson’s isomorphic drawing methods were developed with the aim to revitalise drawing practice in the biological sciences: in the latest body of research, she has researched and developed techniques to widen the conceptual toolkit of scientists and artists studying form understood as in constant flux.

In this talk, Anderson will share specific examples of drawing mitosis in the context of the current AHRC project ‘Representing Biology as Process’, aiming to position drawing as an interdisciplinary tool for the research and visual understanding of biological processes and their dynamic interactions. Anderson will end this section by discussing the epistemological value of drawing in the context of the 'workshop' format. As she will argue, the

workshop can be used as a tool for sharing and evaluating drawing methods and for engaging different collectives with drawing practices. This format can be adapted towards the involvement general audiences for purposes of divulgation and participation, but also between lab team members as a way to collaborate and exchange. The following section explores this latter situation in more detail.

4. “Representing Mitosis (First subproject in the Representing Biological Process project): Outcomes and Interpretations”, (Anderson and Wakefield, 30 mins).

We discuss how drawing techniques have been developed to provide illuminating and informative representations of the active biological process of mitosis, and we reflect on how these techniques can be used as a means of helping to interpret, reflect and theorise mitosis in an interdisciplinary context. We reflect on the experimental process of establishing a series of ‘Drawing Labs’ for the Wakefield Lab and how the question- and process-led nature of these sessions developed insight into significant aspects of Mitosis. As a result of drawing methods specifically developed by Anderson to support direct experiential learning, Anderson and Wakefield co-created a series of images we referred to as ‘Mitosis Pots’.

We discuss the how the main molecular pathways and physical features of mitosis were incorporated and co-created into an outline ‘pot shape’ of mitosis in the embryo of the fruit fly, *Drosophila melanogaster*, the model system that Wakefield has worked with for 20 years and how the features resonated with their co-creators. We then show how these features were extended to mitosis in different model organisms (human tissue culture cell, nematode worm, higher plant) and in dys-regulated (cancer) cells. We discuss the extent to which the shapes and features are intuitive and imaginative, and how much they are how based on quantifiable data. We describe how we tested the generic “true-ness” of the pot shapes by sending the images to a group of mitosis researchers around the world and asking them to identify the ‘kind’ of mitosis occurring. We suggest that if other researchers can recognise it too, then this patterning may be considered as ‘true’ - in the sense of having been organised by mitosis and the organisms, and experimental methods used to get to grips with it.

In this case drawing provides a more engaged access to and reflection on the process of mitosis than merely ‘witnessing’ the mechanical generation of biological representations through various imaging devices. We suggest these exercises may have the potential to influence the methodological and theoretical approach of the science community.

The difficulty of interpreting the images and the potential rewards of doing so point towards further questions and work. Why does the pot have the shape it

does, what are the benefits of having a shape for mitosis? How is it possible to tell the cancerous pot shape apart from the 'normal' pot shape? Is this intuitive? And why is this useful? What is the link between phenomena and image? We are also considering lab experiments for example measuring correlations between energy and aspects of pot shape.

5. Closing discussion chaired by Dupré, 15 minutes.

The function and limit of Galileo's falling bodies thought experiment: Absolute weight, specific weight and the medium's resistance

Rawad El Skaf – Université Paris 1 Panthéon-Sorbonne, France

The epistemological literature on scientific thought experiments (TEs) is mainly built on a-historical analysis of case studies. This is especially lamentable for Galileo's falling bodies because the literature takes it as a canonical example, while the a-historical analysis yields wide disagreements about its conclusion, leading to divergences pertaining to its epistemic function. Thus, leading the epistemological literature on TEs astray and even turning an important debate into a red herring: The Norton/Brown debate revolves, in part, around how Galileo's TE justifies its conclusions; by direct a priori access to laws of nature or by being a deductive argument. Nevertheless, the TE's function is misrepresented as revealing and justifying a law of nature.

The philosophical literature thus needs a more careful historical analysis of Galileo's TE and the following questions answered, before assessing if and how the TE justifies its conclusion(s): What is(are) the conclusion(s) of Galileo's TE? What is its function in Galileo's both argumentative strategies? What is the role of the particulars involved in its scenario? What are the idealisations involved? Are these idealisations justified? Since vacuum could not be explicitly assumed in the TE, then how did Galileo take into account the effects of the medium's resistance? All these questions could be easily answered once we tackle the more general one: Why is the TE's scenario and conclusion restricted to bodies of the same material (i.e. specific weight)?

This paper aims at analysing the function and limit of Galileo's falling bodies TE, which will provide an answer to these questions. First, I retrace the TE's first occurrence in the *De Motu* (1590), where Galileo explicitly states his intention of "seeking the causes of effects". I show that the TE's function is only destructive: it aims at refuting Aristotle's theory of free fall, one of its two principles to be precise, by showing that absolute weight could not cause divergences in speed of free falling bodies of the same material. Second, I

analyse Galileo's both 1590 and 1638 argumentative strategies that followed the same TE, but led him to defend two incompatible theories of free fall. I show that both theories stem from arguments excoagitated by Galileo to analyse specific weight as a causal factor, with conflicting conclusions: in 1590 Galileo argues, following a hasty Archimedean analogy, that specific weight is a causal factor, which leads him to defend that bodies fall, even in void, at different speeds proportionally to their specific weight. While in 1638 Galileo argues, following two arguments and a real experiment on pendulums, that specific weight could not be a causal factor, which leads him to defend that all bodies fall, in void, with equal speeds. Third, I analyse how Galileo ignores an additional effect in his TE, without explicitly assuming vacuum: the medium's resistance frictional effect, proportional to the falling body's surface to absolute weight ratio. This retardation is lesser for larger than smaller bodies, with the same shape. Finally, I conclude by drawing some implications relative to the epistemic debate on TEs.

An Inferential Account of Model Explanation

Wei Fang – Tongji University, China

This essay develops an inferential account of model explanation, based on Mauricio Suárez's inferential conception of scientific representation and Alisa Bokulich's counterfactual account of model explanation. Suárez's inferential account of scientific representation features how competent modelers draw inferences about the target based on the source, wherein the inferences often take the form of transferring over the claims derived from the source onto the claims about the target (Suárez 2004, 2015). Bokulich's basic idea is that a model "explains the explanandum by showing how the elements of the model correctly capture the pattern of counterfactual dependence of the target system" (Bokulich 2011, 39).

Integrating Suárez's with Bokulich's ideas, this essay suggests that the fact that a scientific model can explain is essentially linked to how a modeler uses an established model to make various inferences about the target system on the basis of results derived from the model. The inference practice is understood as a two-step activity: (i) the modeler first entertains the counterfactual structure of the model in various ways such that she can build a whole range of counterfactual statements about the model, and (ii) she then infers from the model to the target by making a range of hypothetical statements that transfer over claims derived from the model onto claims about the target.

It is important to note that it is the modeler who hypothesizes that the counterfactual structure of the model captures its counterpart in the target.

With this in mind, model explanation (ME) then can be understood more precisely as follows:

(ME) It is the modeler who (i) entertains the counterfactual structure of the model by asking Woodward's w-questions, and then (ii) hypothesizes that the claims derived from the counterfactual structure of the model may be applied to its target system.

As shown above, ME consists of two steps. First, since a model can be described as a structure (Weisberg 2013), i.e., a dependence relationship, it follows that variables in the model counterfactually depend on each other. More specifically, changes (or interventions) in explanans variables that figure in the model can be systematically associated with changes in explanandum variables that sometimes take the form of outputs of the model (note that the explanandum variables, represented in the model, are supposed to describe or reproduce their counterparts in the world). As such, the model can be used to answer Woodward's w-questions about itself: we can ask how one variable in the model would change as a result of intervention on another variable in the model. Second, the modeler then hypothesizes that—based on her background knowledge, modeling goals, conceptualization of the target, etc.—the counterfactual dependence relationships derived from the model may be applied to their counterparts in the target. In other words, a kind of inferential relationship can be hypothesized between the model and its target.

To sum up, the inferential account of model explanation holds that model explanation is essentially a two-step activity, in which the first step involves making counterfactual statements about the model and the second involves making hypothetical statements linking the model to the target. In all these steps, the modelers are of paramount importance to the explanation practice.

A ROAMER With a (Wider) View

Luc Faucher – Université du Québec à Montréal, Canada

Psychiatry is in a state of disarray. Its main tool of classification, the DSM, has been under increasing criticism for its lack of validity. The DSM-5 was supposed to remedy the situation, but according to many (Cooper, 2015; Demazeux, 2015), it just failed at it. It is thought that solutions to the situation will have to come from somewhere else. In previous papers, I (with Simon Goyer, 2015 and 2017) have examined the capacity of the Research Domain Criteria initiative (RDoC in the following) to provide a way out of the muddle in which psychiatry is stuck. Though I think the RDoC could make psychiatry move forward, I pointed to some potential problems with the initiative. In this talk, I will engage in some form of comparative social epistemology, that is, I will compare the

features of the process that lead to the establishment of the research priorities figuring in the Strategic Plan for Research of the NIMH (from which the RDoC emanates) and the ones that lead to its European equivalent, the ROAMER (the Roadmap for Mental Health in Europe). In the first part of the paper, I will argue that some of the features of the process that lead to the ROAMER protect it from some of the problems that many (including myself) sees in the RDoC. Indeed, I will show that the process that lead to the ROAMER's proposals is more likely to give rise to what Philip Kitcher has called a Well-Ordered Science (according to him: “[R]esearch would be well-ordered just in case the questions on the agenda are those that would have been selected by representatives of the full diversity of human perspectives ... each of whom was completely committed to addressing the perceived needs of as many people as possible” [2007, 183]).

If I were to stop at the differences in the processes that lead to the establishment of the priorities in research, the explanation of the differences between the two sets of priorities wouldn't be complete. Indeed, there is a crucial difference that explains many features of the priorities identified by the NIMH: it is the fact that it is a “creature” of the NIH. Indeed, while the ROAMER is, so to speak, a “virtual creature” (it is independent of any institution and it merely suggests to the European countries which priorities the should choose); the Strategic Plan of Research is produced by an institute that gets some of the grand lines of its policies from the NIH to which it belongs. So, in the second part of my paper, I will describe the general policies “imposed” by the NIH (mainly the turn from basic science to translational science) to the NIHM and how those had direct impact on the priorities chosen by the latter.

Epistemological reflections on collecting in medicine: What can we learn from the practices of a 19th century Parisian anatomy society?

Juliette Ferry-Danini – Sorbonne Université, France

Medical anatomy and pathology collections are often considered today as minor (and at times inconvenient) parts of history of medicine. The standard narrative of history of medicine paints anatomy and pathology and their museums as having given way to laboratory practices and molecular biology. Modern medicine is thus usually not considered a collecting or museological science. This paper will argue that this narrative is mistaken. In recent years, scholars have argued that collecting is an important way of knowing in science, and an important practice for scientists. Comparisons have been drawn between natural history and data-driven practices in molecular biology

(Strasser 2012a, 2012b) as well as between medical museums and modern biobanks (Tybjerg 2015). The approach of this paper follows in the footsteps of these scholars and is thus twofold: it is first a historical study of the scientific practices of a 19th century anatomy society (the Société d'Anatomie de Paris), with a focus on the collecting practices characterizing it; secondly, it aims to offer some epistemological reflections regarding collecting in medicine and whether links may be drawn between the 19th century and current practices. The Société d'Anatomie de Paris was founded in 1803 then reborn in 1826 under the care of Jean Cruveilhier. I will identify key traits of the collecting practices characterizing the work of the Société, which notably led to the constitution of the Dupuytren medical collection, closed in 2016. Studying such a society is particularly useful because its carefully kept reports give the details of the daily practices of its scientists. A precise reconstruction of why and how collecting relates to knowing in this context can thus be made. As Strasser did for natural history (2012a), I argue that medical collections of the time are analogous to present day molecular data collections in the sense that they didn't only include material objects but historical, contextual and personal data, photography, registries and so on. This paper will also ask whether in the case of medicine, collecting practices do indeed conflict with laboratory practices (as proposed by Tybjerg 2015). In particular, I show that laboratory, the microscope and experimenting, are inextricably bound up with the practices associated with collecting anatomy samples at the time of the Société d'Anatomie de Paris. Finally, in view of connecting with broader questions, I ask what epistemological issues and tensions were already salient at the time and whether those worries have analogues in current medical collecting practices. Comparisons may help thinking about issues in modern medical collecting: for instance, what constitutes a good quality and/or rich medical sample and why is it important?

Strasser, B. J. (2012a). Collecting Nature: Practices, Styles, and Narratives. *Osiris*, 27(1), 303–340.

Strasser, B. J. (2012b). Data-driven sciences: From wonder cabinets to electronic databases. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(1), 85–87.

Tybjerg, K. (2015). From Bottled Babies to Biobanks: Medical Collections in the Twenty-First Century. In *Fate of Anatomical Collections* (Rina Knoeff; Robert Zwijnenberg., pp. 263–278). Farnham: Ashgate.

Innovation by incomplete theorization: The case of direct cell reprogramming

Grant Fisher – Korea Advanced Institute of Science and Technology, Republic of Korea

This talk contributes to philosophies of scientific practice by exploring a case of an emerging biotechnology in a complex ethical, legal, policy, and commercial context. Direct cell reprogramming is aimed at generating pluripotent stem cell lines for basic and therapeutic research while avoiding the destruction of human embryos. As soon as it was introduced in the last decade, direct cell reprogramming came to be regarded as supposedly providing a “scientific solution to an ethical dilemma” (Rao & Condic, 2008). In response, some of the bioethics literature points its failure as a “technical solution” (Devolder 2015). However, interest in direct cell reprogramming was garnered among researchers whose normative strategies were not merely aimed at avoiding disputes over the ethics of embryo destruction. These strategies were also aimed at gaining access to technologies capable of enabling further research in stem cell biology. In this respect, human embryonic stems remain indispensable to stem cell biology (Fagan 2013). But policy and commercial restrictions on the derivation and use of human embryonic stem cell lines in the United States may also have had negative epistemic impacts on stem cell research. Direct cell reprogramming came to prominence during a period of policy restrictions on federal funding for human embryonic stem cell research subsumed under a complex and heterogeneous policy framework, amid disputes over the interpretation of relevant law, confusion over how to interpret policy under Presidential administrative change, and a highly competitive and restrictive patent culture. This talk explores normative strategies underlying the development of direct cell reprogramming – strategies that seek agreements while attempting to be circumspect with regards to divisive fundamental “theoretical” issues. One well known example of such a strategy – albeit originating in a different context – is what Cass Sunstein (1996) calls an “incompletely theorized agreement”. In its most basic sense, an incompletely theorized agreement is a legal decision reached without getting into theoretical issues that may divide political opinion. Sunstein argues decision makers ought to refrain from engagement with these divisive issues in the interests of stability, mutual respect, and reciprocity. The present paper aims to employ an analogous sense of “incomplete theorization” in a way that is useful to understanding the governance of emerging biomedical science and technologies. It will be argued that while direct cell reprogramming provided no technical solution, it is a case of an attempted innovation by incomplete

theorization because it proposed to allow researchers to pursue a range of normative aims clustered around epistemic, therapeutic, and economic values while seeking to refrain from engagement with fundamental legal, political, and ethical disputes. The factual failure regarding some of the practitioners' aims is only part of the issue. The normative problem concerns the legitimacy of incomplete theorization as a means to forge agreements over the direction of scientific research in pluralistic societies.

Devolder, K. *The Ethics of Embryonic Stem Cell Research* (Oxford: Oxford University Press, 2015).

Fagan, M.B. *The Philosophy of Stem Cell Biology: Knowledge in Flesh and Blood* (Basingstoke: Palgrave MacMillan, 2013).

Rao, M. & Condic, M.L. "Alternative Sources of Pluripotent Stem Cells: Scientific Solutions to an Ethical Dilemma", *Stem Cells and Development*, 17: 1-10.

Sunstein, C.R. *Legal Reasoning and Political Conflict* (Oxford: Oxford University Press, 1996).

Brain metaphors and dementia scientific discourse

Giulia Frezza – Sapienza University of Rome, Italy

This paper presents a study of philosophy of science in practice in the medical field. Metaphor is dominant in dementia public discourse both at individual and societal level (cf. Alzheimer Europe's 2013). Dementia is metaphorically understood as a 'tidal wave' threatening our society, but it is also a disease metaphorically turning people into 'vegetables' or 'zombies' in the eyes of their family and caregivers.

The role of metaphor in dementia discourse is crucial, as metaphor is commonly used to communicate complex issues in simpler terms. However, as shown lately by the literature, metaphor use varies among language users and it may lead to misunderstanding and social stigma (Lane et al. 2013; Zeilig 2014).

In our study, we analyzed 205 articles in health communication about dementia by means of Wmatrix (a software tool for corpus analysis and comparison). We showed how the two opposed metaphors of 'decline' vs. 'plastic' brain frame two opposite narratives, defined by different semantic domains and involving opposed ideas of time. The first, decline, is embedded within the domains of 'curing' and 'death-related problems', while the latter, plasticity, underlines brain's power, by means of individual's resilience, training and learning capabilities. Moreover, while the narrative of decline emphasizes

the category of time as ‘passing and elapsing’, the narrative of plasticity defines a view of time as ‘forming’, ‘starting’ and ‘renewing’, focusing on activity and revealing a counter-time where ‘it’s never too late’.

For better understanding where the current situation comes from –also in view of the peak of dementia ‘epidemic’– we will frame it within a historical-epistemological background, discussing the change of diagnosis, nosology and therapy models of dementia through times.

Scientists’ language about dementia has indeed changed through times, culture and scientific paradigms: in XIX century, psychiatrists such as Pinel, used dementia to refer to patients with ‘intellectual deterioration and idiocy’, and Kraepelin distinguished a ‘terminal dementia’ and a ‘general decay of mental efficiency’, both focusing on decline; recent brain studies, by contrast, emphasize the notions of ‘brain plasticity’ and ‘cognitive reserve’, which highlight a different individual susceptibility to age-related brain changes.

In conclusion, the analysis of metaphor use in dementia underlines what we define different ‘ethical risks and responsibilities’ in health communication, that is the relation between the effective risks and the responsibilities that different targets of people (as physicians and journalists) assume in communicating metaphors to different targets of audience (as patients, family, caregivers and society at large). Responsible metaphor use in dementia discourse exhibits important implications for fine-tuning communication strategies in health-care and therapy (like spreading information about how slowing down the disease, especially in the window period) and for fostering positive health behaviour.

Alzheimer Europe Report-2013:

[http://www.alzheimer-europe.org/Ethics/Ethical-issues-in-practice/2013-The-ethical-issues-linked-to-the-perceptions-and-portrayal-of-dementia-and-people-with-dementia/How-dementia-is-perceived-and-portrayed/Metaphors/\(language\)/eng-GB#fragment4](http://www.alzheimer-europe.org/Ethics/Ethical-issues-in-practice/2013-The-ethical-issues-linked-to-the-perceptions-and-portrayal-of-dementia-and-people-with-dementia/How-dementia-is-perceived-and-portrayed/Metaphors/(language)/eng-GB#fragment4);

Lane H. P., Mclachlan S., Philip J., The war against dementia: are we battle weary yet? *Age and Ageing* 2013; 42: 281–283;

Zeilig, H. (2014). Dementia as a cultural metaphor. *The Gerontologist*, 54, 258–267.

Causal Selection Problems in Epidemiology: The Case of Increased Incidence of Psychosis in Ethnic Minorities

Katherine Furman – University College Cork, Ireland

Hannah Jongsma – University of Cambridge, United Kingdom

Research in psychiatric epidemiology indicates that rates of psychosis are higher amongst ethnic minorities in Western European countries than the White majorities in those same countries. Current best research suggests that any plausible account of this finding will be multi-causal, potentially including (amongst other factors): vitamin-D deficiency; stressors of migration; social exclusion; ‘psychosocial disempowerment’ (lacking control over one’s own life); economic precariousness, etc.

In some ways, this result can be easily accommodated within current causal thinking in epidemiology, due to the increased acceptance of multi-causal theories of disease since the mid-twentieth century. This is especially clear in the adoption of Rothman’s (1976) ‘causal cakes’ into epidemiological thought, which will be familiar to philosophers as closely resembling Mackie’s (1965) INUS conditions (“Insufficient but Necessary parts of a condition which is itself Unnecessary but Sufficient”).

The metaphor of the causal cake is helpful for admitting multiple causes to the explanation, but unhelpful in that it fails to offer a way of distinguishing which causal factors are more or less important to the effect. Being able to distinguish the relative importance of causes is pragmatically important, both for the treatment of individual patients (where should the intervention be targeted?) and for developing policy interventions (given resource constraints on policy makers, policies should be targeted at the most salient causes).

Susser’s (1973) nested system model of disease introduced the idea of ranking the relative importance of causes in epidemiological thought by suggesting that the causes most proximal to the effect be considered the most salient. However, Susser’s account relies on a dubious metaphysics of levels, and does not allow for distant causes to ever be the most salient. For instance, we might think that smoking as the most salient cause of lung cancer even though degradation of lung tissue is the most proximal cause.

In this paper we argue that the problem of assessing the relative importance of causes in epidemiology is recognisably a causal selection problem from Philosophy of Science, and should be treated as such. In particular, we take seriously Woodward’s (2010) suggestion that in assessing the relative salience of causal factors one should look at how specific a causal factor is to its effect, and how stable it is when surrounding causal factors are varied. However, we also argue that the case of increased psychosis in ethnic minorities (the central

case of this paper) requires that more careful attention be paid to what it means for a cause to be 'stable'. In particular, Jongsma (2017) has argued that psychosocial disempowerment is the most salient cause of increased rates of psychosis amongst ethnic minorities. However, how would one separate this cause from other closely related factors, such as social exclusion and economic precariousness, in order to test its stability? In this paper, we make some suggestions for how to adapt Woodward's account of causal selection to make it more practically useful for epidemiologists dealing with cases such as ours.

On the Construction of Scientific Narratives

Devin Gouvêa – University of Chicago, United States

In 1998, Frederick Suppe presented a painstaking microanalysis of a seminal scientific paper as evidence that several leading models of testing and confirmation were hopelessly misguided. Critics took issue with almost every aspect of Suppe's analysis and conclusions, but none challenged his decision "to treat the scientific paper as the vehicle of testing and confirmation" (384). Suppe grounded this move by an appeal to the positivist conception of scientific justification as publicly accessible and clearly separated from the more inscrutable processes of discovery.

In the intervening decades, philosophers and historians of science have challenged the central distinction between discovery and justification and have paid increasing attention to other aspects of scientific practice besides the publication of results. As Schickore (2008) has shown, the analysis of Suppe and colleagues is part of a diverse cross-disciplinary response to the widely recognized mismatch between the practices of scientific investigation and the presentation of its results. Not everyone agrees that this mismatch is epistemically significant, let alone what its significance may be.

In this talk, I revisit that mismatch from two complementary perspectives — that of a scientist with experience instantiating it and a philosopher convinced that it is epistemically significant. I argue that the process of constructing scientific narratives is not merely incidental to the establishment of reliable scientific claims. Instead, this process, though susceptible to distortion, has a positive role to play in the testing and confirmation of scientific knowledge.

Using my own personal archive of lab notebooks, paper drafts, and meeting notes, I identify three particular manifestations of the mismatch. First, the main characters of our narrative — particular experimental phenomena and their putative explanations — were not established in advance, but iteratively refined during the course of the study. Second, there was a complex, non-linear mapping between the order of experimentation and the logic of the

narrative. The experiments that occupied the most time and effort did not necessarily turn out to be the most central to the published narrative. Third, the narrative nevertheless uses data very efficiently. Data from a single fruitful experiment is generally distributed throughout the narrative, while less productive experiments are retained in minor supporting roles.

Though our project was motivated by clear biological questions, these were not sufficient to determine our path through a series of methodological crossroads. We had more experimental questions that we could possibly pursue, more patterns in our experimental data than we could possibly explain, and more ways of representing any individual experiment than a paper could possibly hold. Our choices were conditioned by pragmatic considerations but ultimately governed by epistemic norms.

The questions we faced are pertinent to many kinds of experimental projects. Given the data within reach, which of the conceivable experimental phenomena are most stable, most biologically meaningful, and most susceptible to explanation? At its best, the constraints of a published narrative encourage scientists to produce their strongest arguments. At worst, they create incentives and pressures that risk undercutting good epistemic practice. Either way, we cannot escape narrative constraints, and we should not want to.

Gouvea et al. (2015), Experience Modulates the Reproductive Response to Heat Stress in *C. elegans*, PLOS One 10

Schickore (2008), Doing Science, Writing Science, Philos Sci 75

Suppe (1998), The Structure of a Scientific Paper, Philos Sci 65

Risky Measurements: Scientific Strategies of Dealing with Instability and Divergence in the Measurement of Risk Preference

Till Grüne-Yanoff – Royal Institute of Technology (KTH), Sweden

The concept of risk preferences is central in economics and psychology. It is used for descriptive and explanatory purposes: to predict behavior under changes in risk, and to explain social phenomena involving uncertainty. It is also used for normative purposes: to determine how a possible change in uncertainties would affect individuals' welfare. For these purposes, behavioral scientists have conceptualized risk preferences as context-, domain- and sometimes even individual-independent. Yet the now long-standing practices of measuring such a concept have yielded multiple divergences and instabilities that provided at best mixed evidence for such independencies. In this paper, I

investigate the different strategies that scientists have adopted in the face of such adverse measurement results.

Initial attempts to measure risk preferences revealed considerable heterogeneity between individuals. Attempts to explain these differences through demographic factors have largely failed (Eckel and Grossman 2003). Instead, focus then shifted to explain heterogeneity as genuine individual idiosyncrasies. This raised the issue of correct measurement methods and constructs, leading to numerous methodological investigations. Three main methods can be distinguished: (i) behavioral tasks in the laboratory, (ii) questionnaire-based surveys, either with abstract questions or hypothetical scenarios and (iii) self-reported actual risk activities. Each of these methods has produced many constructs, which also reveals a conceptual uncertainty about the nature of risk preferences. For example, behavioral constructs include the choice list, ranking and allocation procedures, survey constructs include the gambling attitude and beliefs survey and the personal risk inventory, and actual behavior constructs include the alcohol use disorders identification test and the encounters with risky situations survey. Although there are some convergence results for survey constructs (Dohmen et al 2011) and between survey and actual behavior constructs (Lönngqvist et al. 2015, Frey et al. 2017), these results are still weak and do not answer worries about lack of divergence validity. Furthermore, the possibility of method factors – i.e. that convergence is largely produced by the nature of the measurement method, rather than by the similarity of constructs – remains not investigated. And finally, even if these worries could be answered, this approach still must deal with substantial evidence of within-person variability (Bardsley 2009, ch. 7).

Researchers have developed four different strategies to deal with these divergence problems. First, the standard response is to stick to the true score theory, assuming a deterministic core plus white noise, and potentially making refining assumptions about the error term (Blavatzkyy & Pogrebna 2010). The second is to propose a bifactor model, assuming a general factor of risk preference, combined with a battery of specific factors (Frey et al. 2017, 6). A third strategy is to revert to constructed preference theory, and consider the measurement a result of a deliberative process in which the agent relies on a set of possible preferences to construct a response (Loomes & Pogrebna 2014). A final option is to give up on a global concept of risk preference altogether (Berg et al. 2005, Berseghyan et al. 2011).

Comparing these four strategies, I show that each of them requires specific theoretical and conceptual adjustments, as well as modifications of the respective purposes in which the constructs thus stabilized can be applied. Instead of prescribing a single strategy, I am thus more interested in the respective coherence of each strategy, and what particular conceptual and practical implications they entail.

Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C. and Sugden, R. 2009. *Experimental Economics: Rethinking the Rules*, Princeton, NJ: Princeton University Press.

Berg, Joyce, John Dickhaut, and Kevin McCabe. 2005. Risk preference instability across institutions: A dilemma. *Proceedings of the National Academy of Sciences of the U.S.A.* Vol. 102 No. 11: 4209-4214.

Berseghyan, Levon, Jeffrey Prince, and Joshua C. Teitelbaum. 2011. Are Risk Preferences Stable across Contexts? Evidence from Insurance Data. *The American Economic Review* Vol. 101 (April): 591-631.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., & Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3), 522-550.

Eckel, Catherine C. and Phillip J. Grossman, —Sex and Risk: Experimental Evidence. 2003 in C. R. Plott and Vernon L. Smith, eds. *Handbook of Results in Experimental Economics*. Amsterdam: North-Holland/Elsevier Press.

Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. 2017. Risk preference shares the psychometric structure of major psychological traits. *Science advances*, 3(10), e1701381.

Lönnqvist, J. E., Verkasalo, M., Walkowitz, G., & Wichardt, P. C. 2015. Measuring individual risk attitudes in the lab: Task or ask? An empirical comparison. *Journal of Economic Behavior & Organization*, 119, 254-266.

Colorful Boxes

Reiner Hähnle – Technical University of Darmstadt, Germany

A black box is an artifact whose inner workings and construction principles one does not (fully) understand, yet one expects to function nonetheless, and even relies on it, perhaps with one's life. It is a stance we got used to taking towards the technical devices of ever increasing complexity we surround ourselves with. In fact, the black box principle is frequently a metaphor for the impossibility to understand the complexity of human-created technology. Even worse, to merely open up a black box is insufficient, because its implementation is described with the help of further black boxes which contain black boxes in turn, and so on. Thus the black box principle is associated with a loss of the ability to fully comprehend complex technological devices, in particular, when these are composed of software. It is, unsurprisingly, often found in technology-skeptical discourses.

For this reason, it is important to highlight the existence of a second interpretation of the black box principle that serves to reduce descriptive

complexity. It consists of two parts: first, a precise description of the boundaries and the expected behavior of a black box that we call its interface. Second, a rigorous argumentation that the content of a black box faithfully implements this interface. Assume we encounter a complex artifact: The first step allows us to replace it by its interface, essentially treating it as a black box. The second step tells us we do not lose information in doing so. Whenever an interface is easier to characterize than its realization, replacing the latter by the former results in a potentially drastic reduction of descriptive complexity. It also mitigates the problem of "nested boxes".

We illustrate this idea with a simple computer program, an algorithm that sorts a given list of names in ascending order. The interface of this particular black box is easy to describe: "Input is a list of names, output is exactly the same list of names arranged in ascending order, and there is no other effect." The last clause is needed to exclude any unwanted interference of our black box with its environment. Sorting algorithms can be highly complex: even wide-spread implementations exhibit bugs. However, it is possible to formally prove that an algorithm implements its interface. Hence, to make precise use of the black box "sorting", it is unnecessary to access an actual sorting algorithm: its vastly simpler interface suffices.

In Computer Science, the principle to distinguish between an interface and its implementation, with the intention to master complexity, goes back to the late 1960s, to the inception of object-oriented programming and abstract data types. Even earlier, the systematic investigation to relate an interface in the sense of a logical theory to the mathematical structures that realize it, was the subject of model theory since the 1950s.

Obviously, the interface/implementation construction of black boxes has limitations. For example, some machine learning algorithms seem to defy the construction of interfaces that are substantially less complex than their implementation. But it is not impossible in all cases and it is important to understand where the principal limitations lie. We argue that black boxes need not merely be black. They should not primarily be seen as a metaphor for the loss of comprehensibility in the presence of overwhelming technical complexity, but as a structuring principle to master it.

Epistemic and methodological dimensions in interviews with scientists talking about challenges in their day-to-day research practice

Nora Hangel – Indiana University Bloomington, United States

Scientists find themselves increasingly challenged to articulate and defend their strategies and decisions to broader publics. However, studies of scientific practice have only started to consider the many socially relevant dynamics of knowledge production in relation to methodologically relevant challenges. Especially in collaborative and interdisciplinary environments, scientists' abilities to give clear accounts of the methodological norms and standards are a crucial prerequisite for the functioning of scientific practice.

How scientists themselves conceptualize situations of choice in non-algorithmic decisions and selection criteria in scientific inquiry has not been systematically analyzed. To analyze these aspects of choice and decision-making in knowledge production, we [1] use interviews with scientists from a broad range of disciplines in the natural and social sciences, including medical sciences and engineering

For an empirically informed POS the internal epistemic dimension stays central to the investigation and all other aspects are integrated in respect to their epistemic and methodological relevance. When we analyzed narrative reconstructions of scientists talking about their practice, we did not observe what scientists actually do. Thus, the reflections, experiences, and challenges scientists shared in the interviews are empirically grounded but nevertheless contribute to a conceptual framework about how science is done.

How did we proceed? On the one hand, we drew on philosophical debates about theory choice, exploratory research, epistemological strategies for experimentation, and values in science as guides for the exploration of scientists' reflections about their judgments in scientific inquiry. On the other hand, we use qualitative data analysis to identify patterns emerging from comparing 'quotes' – or semantical units – of scientists when they speak about their day-to-day experiences and challenges.

When analyzing concrete accounts expressed in examples of their own experience, we found numerous expressions of uncertainties not yet addressed and integrated in the POS discourse. We found that while scientists routinely express broader methodological commitments, they regularly express their doubts about how criteria such as reproducibility or fit with accepted theories should be put to work in concrete situations. We also found accounts that scientists distribute the responsibility for ensuring the quality of research results to others (e.g., collaborators, referees, also replication efforts

are delegated to others). On the other hand, referees confess to have limited resources for checking the newly presented results, so they rely on the researchers' epistemic ethos for doing sound research. Furthermore, we found a conflation of error, fraud and technical or personal failure as well as a mismatch of applying norms and standards when promoting their own research and evaluating research of others. By analyzing how scientist conceptualize methodologically and epistemically relevant challenges in their research practice we combine qualitative research methods and discourse analysis. Thus, we enhance the traditional analytic framework with empirically grounded data, and thereby seek to integrate social and pragmatic factors in relation to methodological issues.

Hangel, N. & Schickore, J. (2017) "Scientists' conceptions of good research practice" *Perspectives on science*, Volume 25, Issue 6. doi: 10.1162/POSC_a_00265.

Hangel, N. & Schickore, J. (submitted) "Scientists' views on choices and decision-making in scientific practice".

Challenging Van Fraassen's observable-unobservable distinction: a case study from nineteenth century medical microscopy

Mahi Hardalupas – University of Pittsburgh, United States

An important issue dividing constructive empiricists and realists is the role of observation in the production of scientific knowledge and the observable-unobservable distinction. Realists, such as Maxwell (1962), argue for a 'continuum of vision' where there is no fundamental distinction between observable and unobservable entities because it is possible to extend the senses through advances in scientific instruments and technology. In developing constructive empiricism, Van Fraassen (1980, 2008) revives a version of the observable-unobservable distinction claiming that "X is observable if there are circumstances which are such that, if X is present to us under those circumstances, then we observe it." (1980, 16) For Van Fraassen, observation is unaided perception and thus what is observable is bounded by the limits of the human senses. Whether one "observes" through a microscope is a central example in this debate, (Hacking, 1985; Kusch, 2015) where realists claim that you do observe with a microscope and constructive empiricists disagreeing. However, despite the importance of the microscope as an example, little work has looked to scientific practice to understand how observation featured in historical debates surrounding the use of the microscope.

In this essay, I examine the history of microscopy in nineteenth century Britain, specifically focusing on the epistemic attitudes of medical researchers towards the use of high-power microscopes (ie. microscopes with high magnifying powers) as an instrument for observation. By integrating a variety of primary sources, I show that in the mid-nineteenth century, physicians and natural historians such as Darwin distrusted high-power microscopes and favoured low-power microscopes that closely corresponded to what was visible with the unaided senses. Prima facie, this appears to support Van Fraassen's distinction and its reliance on the unaided senses. However, a more careful examination of nineteenth century microscopic practices counters this initial interpretation. While it is true that the unaided senses were considered reliable generators of knowledge, I explain how high-power microscopes were typically used in tandem with low-power microscopes, which acted as mediators between high-power microscopes and the unaided senses. This demonstrates that, contrary to Van Fraassen, early medical microscopists embraced a continuum of vision where the senses could be extended through instruments. Furthermore, I show there is evidence that nineteenth century physiologists understood direct observation as including microscopic observation.

I conclude that this historical case study challenges Van Fraassen's unobservable-observable distinction. This need not render constructive empiricism implausible but does show that Van Fraassen's observable-unobservable distinction conflicts with historical scientific practice in microscopy.

Hacking, I. (1985). Do We See through a Microscope? *Images of Science: Essays on Realism and Empiricism*.

Kusch, M. (2015). Microscopes and the Theory-Ladenness of Experience in Bas van Fraassen's Recent Work. *Journal for General Philosophy of Science*, 46(1), 167–182.

Maxwell, G. (1962). The Ontological Status of Theoretical Entities. *Minnesota Studies in the Philosophy of Science*, 3–27.

van Fraassen, B. C. (1980). *The scientific image*. Oxford: Oxford University Press.

van Fraassen, B. C. (2008). *Scientific representation: Paradoxes of perspective*. Oxford: Clarendon Press.

A patchwork approach to endogenous brain activity in resting state neuroimaging research

Philipp Haueis – Humboldt University Berlin, Germany

Resting state functional connectivity studies are widely used in neuroimaging, but the exact functional roles of endogenous brain activity traced by this methodology remains debated. Some argue that it plays a direct role in cognition and behavior (Margulies et al. 2016), while others suggest that it enables information processing by maintaining functional brain systems (Raichle 2015). Despite this debate among practitioners, philosophers have exclusively focused on direct cognitive roles when describing how resting state neuroimaging can be used to revise cognitive architecture (Bechtel 2013), characterize long-term mental processes (Klein 2014) or draw psychological inferences (McCaffrey and Danks in press). These accounts are incomplete because functional roles can vary depending on type and frequency of the resting state signal (Keilholz et al. 2017).

In this paper, I overcome this incompleteness by proposing a patchwork model which describes how neuroimaging researchers can measure different functional roles of endogenous brain activity. The patchwork model assumes that scientific concepts like “functional connectivity” have different local applications depending on the neuroimaging technique used to measure endogenous brain activity (cf. Wilson 2006). I distinguish three functional roles by their causal specificity in cognition and behavior (Woodward 2010). Activity that is operative in cognition is causally specific: manipulating it has a graded and task-specific effect on behavior. Modulatory activity has a graded but task-unspecific effect on many behaviors. Manipulating activity that enables cognition has nonspecific effects that switches many behaviors “on” or “off”.

The resulting patchwork model shows how four different neuroimaging techniques can be used to measure endogenous brain activity with operative, modulatory or enabling functional roles. Combined positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) studies measure the regional variation of metabolic resources (Vaishnavi et al. 2010). Manipulating resources has causally nonspecific effects because they enable multiple neural processes (neuron repair, synaptic learning). fMRI studies using a quasi-periodic pattern algorithm can measure modulatory activity because it tracks changes in physiological arousal (Thompson et al. 2014). Such changes affect many behaviors in a graded fashion. Finally, fMRI, electroencephalography (EEG) and electrophysiology techniques can track endogenous activity that is operative in cognition and behavior. Such studies show task-specific changes in endogenous activity during motor tasks (Fox et

al. 2007), frequency-specific EEG signal changes during cognitive tasks like reading (Palva and Palva 2012), or temporally specific endogenous firing patterns that reflect the statistical history of sensory inputs (Schölvinck et al. 2013).

The patchwork model overcomes the bias of existing philosophical accounts towards operative conditions because it reveals that current neuroimaging techniques can be used to measure a whole variety of endogenous brain activity. It also provides a conceptual resource that resting state researchers can use, revise or extend when they characterize endogenous brain activity with enabling, modulatory or operative roles in cognition and behavior.

Bechtel, W. (2013). The Endogenously Active Brain: The Need for an Alternative Cognitive Architecture. *Philosophia Scientiae* 17, 3–30.

Fox, M.D. et al. (2007). Intrinsic Fluctuations within Cortical Systems Account for Intertrial Variability in Human Behavior. *Neuron* 56, 171–184.

Keilholz S.D. et al. (2017). Noise and Non-Neuronal Contributions to the BOLD signal: Applications to and Insights from Animal Studies. *Neuroimage* 154, 267–81.

Klein C (2014) The brain at rest: what it's doing and why that matters. *Philos. Sci.* 81(5), 974–85.

Margulies D.S. et al. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization *PNAS* 113(44), 12574–79.

McCaffrey J., Danks, D, (in press). Mixtures and psychological inference with resting state fMRI. *Brit. J. Philos. Sci.*

Palva J.M. et al. (2012). Infra-Slow Fluctuations in Electrophysiological Recordings, Blood-Oxygenation-Level-Dependent Signals, and Psychophysical Time Series. *Neuroimage* 62(4), 2201–11.

Raichle, M. (2015) The Restless Brain. How Intrinsic Activity Organizes Brain Function. *Trans. Roy. Soc. Lond. B* 370(1668), 1–8

doi: 10.1098/rstb.2014.0172.

Schölvinck M.L., et al. (2013). The Contribution of Electrophysiology to Functional Connectivity Mapping. *Neuroimage* 80, 297–306.

Thompson G.J. et al. (2014). Quasi-Periodic Patterns (QPP): Large-Scale Dynamics in Resting State fMRI that Correlate with Local Intraslow Electrical Activity. *Neuroimage* 84, 1018–31.

Thermophysical Reference Data as Memory Practice: Scientific Facts and Techniques of Forgetting

Elliott Hauser – University of North Carolina at Chapel Hill, United States

In this paper I examine the production of thermophysical reference data as a memory practice, defined as a material instance of constructive forgetting (Bowker 2005). I'll show how the manifold material traces produced via experiment are reduced into formal inscriptions. These inscriptions, promulgated as reference data, enable the scientific forgetting of the underlying empirical results. Bowker's insight was that this process is jussive, and I contend that studying the archive's resultant exclusionary power can yield deep insights into the relationship between inscription practices and scientific knowledge.

Thermophysical reference data are an ideal site for an exploration of the processes and techniques of forgetting. As a specific example, take reference data for the thermal conductivity of copper (Ho, Powell, and Liley 1974). Figure 1a) combines a comprehensive review of empirical results alongside a Recommended line. Figure 1b) shows an abridged version containing just the Recommended line, intended for easy direct reference and incorporation into scientific handbooks. In a recent version of the CRC Handbook of Chemistry and Physics (Haynes 2012), the thermal conductivity of copper is presented as a table, the values more or less adhering to Ho et al.'s recommendations. The erasure of detail and context by this reification into reference data forms a material record of how this specific memory practice enacts Bowker's forgetting.

The fact that Ho et al.'s graphs and tables constitute material artifacts of forgetting may seem like a challenge to their work's rigor. On the contrary, it is precisely their rigor which scientifically licenses the forgetting that has subsequently ossified the fact of copper's thermal conductivity. However, the suggestively empirical nature of the term 'reference data' is indeed problematic, especially in light of the erasure of empirical results just described. This is 'data' only in the sense of 'a given'. More broadly, I argue that reference data artifacts are precisely designed to minimize what Star has called "residual categories" (Star 2010). Thermophysical facts resemble an 'ideal type' boundary object (Star and Griesemer 1989, 410–411), bridging thermophysical properties research with its military and engineering applications. However, this bridge seems to be far more permanent, and less negotiated, than those found in the sciences Star studied, like paleontology. The sufficiency of thermophysical facts within aerospace engineering, the primary consumer (and funder) of thermophysical

Star, Susan Leigh, and James R. Griesemer. 1989. "Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39." *Social Studies of Science* 19 (3). SAGE Publications Ltd:387–420.

The Species Category as a Scientific Kind

Caleb Hazelwood – Georgia State University, United States

Marc Ereshefsky's project of eliminative pluralism is simply stated in two theses: 1) In light of the myriad mechanisms of speciation legitimated by scientific practice, we ought to be pluralistic realists about species taxa, and 2) as there is no unifying feature among all species taxa, we ought to doubt the existence of the species category. Here, I will argue that one promising strategy for saving the species category is to conceive of it as a natural kind after the practice turn. I will do this by situating the species category within a recent practice-based account of natural kinds proposed by Marc Ereshefsky and Thomas Reydon called "scientific kinds." Scientific kinds are legitimate natural kinds. They enforce ontological boundaries, not merely epistemic ones. Most importantly, they recognize boundaries drawn from the lab and the field, not only from the armchair.

According to an account of scientific kinds, the species category is perfectly real by virtue of the same principles that legitimize various species concepts: it is a category determined by the epistemic aims, methodologies, and classificatory practices of our best science. The species category does theoretical and explanatory work in scientific practice, so a practice-based theory of kinds has reason to legitimize it. In a recent paper, Adrian Currie demonstrates this point with a case study of scientific practice. Currie mounts a defense of the species category based on the indifference to species concepts in paleobiology. When establishing new species, paleobiologists use a set of criteria that are entirely indifferent to the specifications that delimit one species concept from another. That is to say, their explanatory pursuits range across a myriad of species concepts. No one species concept motivates taxonomic practices in paleobiology; instead, the species category itself does a significant amount of explanatory heavy lifting.

This example is of course insufficient to provide a full defense on behalf of the species category as a scientific kind. It does, however, provide some insight into how the taxonomic rank may be more than a conceptual vestige. If these insights into paleobiological practice are representative of the species category's theoretical utility, even if only in some disciplines, then it is demonstrably progressive when compared to the eliminativist approach. This

does not amount to a full-throated defense of scientific kinds, either, as the account is not without problems. The point of this exercise is to situate the species category within an account of natural kinds that is largely sensitive to scientific practice instead of pure theoretical coherence. This, I argue, will be necessary to save the species category: instead of ignoring its heterogeneity in pursuit of some metaphysically unifying feature, we ought to embrace its internal differences and recognize that they do not preclude its classificatory power in scientific practice.

Phylogenetic Taxonomy and HGT: Reconcilable Tension or Grounds for Taxonomic Revolution?

Phillip Honenberger – University of Pittsburgh, United States

Phylogenetic taxonomy, in the sense articulated by Hennig (1966) and since defended and developed by many others (e.g. Wiley 1981, Wiley and Lieberman 2011), proposes to classify organisms on the basis of nested separations of lineages in the course of evolutionary history. Yet the increasingly appreciated frequency of horizontal gene transfer – that is, cases wherein an organism receives genetic material from the environment or other organisms by non-reproductive processes, such as viral insertion – has led many to challenge key features of the phylogenetic perspective (e.g. Doolittle 1999, 2010). (HGT is most common in microbial populations but other forms of introgression, such as endosymbiosis and hybridization, present similar problems for phylogenetic taxonomy.) A major question is whether and to what extent phylogenetic taxonomy is an appropriate approach for organisms, species, and lineages wherein HGT has played a role.

In the first part of the talk, I seek to articulate the tension between HGT and phylogenetic taxonomy as clearly as possible. I begin with a review of prominent phylogenetic taxonomists' statements of its main aims (from the 1960s to today) and of their motives for recommending it as the basic framework for biological classification. It is then shown how and why widespread HGT would make the realization of some of these aims impossible, as well as weakening the justification for preferring phylogenetic taxonomy to alternative taxonomic frameworks. The main difficulties arise from (a) the insufficiency of bifurcating tree models to capture all major evolutionary-historical events relevant to speciation and character evolution in a HGT-prevalent population, as well as (b) the ambiguousness of hierarchies constructed on the basis of tree models that include reticulation.

Assuming that phylogenetic taxonomists will want to respond to these challenges, I spend the second half of the talk describing and evaluating some

possible strategies of resolution. These strategies are chosen as ones that preserve as much as possible of the central motivations and the accumulated tools of phylogenetic taxonomy without ignoring or denying the phenomenon of HGT. In short, I examine options for “preserving” phylogenetic taxonomy as the basic taxonomic framework for biology, despite the challenges of HGT. These solutions involve a mixture of theoretical, practical, and modeling/representational components (e.g. tree diagrams of various sorts). Some of the options here include (i) deny the applicability of phylogenetic systematics to introgressed lineages [“quarantine” the problem]; (ii) found phylogenetic systematics on “trees of cells” rather than character trees or gene trees; or (iii) incorporate reticulation into phylogenetic models (via representations of: speciation-by-hybridization, transfer events, mosaic inheritance, or overlapping trees) [e.g. Huson & Bryant 2005, Kunin et al. 2005, Baptiste and Burian 2010].

These strategies of reconciliation are not entirely unfeasible, even if none evades at least some of the troubling implications for phylogenetic taxonomy detailed in the first part of the talk. However, the third set of strategies especially ((iii) above) may enable construction of a taxonomic framework that preserves major aims and motives of phylogenetic taxonomy while incorporating HGT-related phenomena into taxonomic representation and reasoning.

Analogical Reasoning: Lessons from Davy’s Work on Electrochemical Decomposition

Jonathon Hricko – National Yang-Ming University, Taiwan

Yafeng Shan – Durham University, United Kingdom

In his textbook *Elements of Chemical Philosophy* (1812), Humphry Davy regards analogy, along with observation and experiment, as the three fundamental methods that chemical philosophers use to acquire knowledge about the world. Moreover, analogy plays a particularly important role. As Davy puts it, “in the progression of knowledge, observation, guided by analogy, leads to experiment, and analogy, confirmed by experiment, becomes scientific truth” (1812, p. 1). Our goal in this paper is to examine the nature of analogical reasoning in Davy’s work, and to explore its consequences for our understanding of analogical reasoning in scientific practice.

We focus on the work in electrochemistry that Davy presented in his 1806 and 1807 Bakerian Lectures. Central to this work is Davy’s use of the Voltaic pile to decompose various compound substances. Davy began by decomposing water, and then moved on to decompose other compound substances. Eventually, he

succeeded in decomposing two previously undecomposed substances, namely, potash and soda, and he thereby discovered potassium and sodium.

We identify three roles of analogical reasoning in Davy's work that illustrate the ideas regarding analogy presented in his textbook. First, Davy uses analogies to infer more general hypotheses. After observing that electricity has the effect of decomposing a particular substance, Davy hypothesizes that electricity may also decompose other, analogous substances. Second, analogies guided Davy in the design of the experiments used to test these hypotheses. Davy's use of the Voltaic pile to decompose various substances depended on modifying the initial experimental setup that he used to decompose water. These modifications were often motivated by analogies between the substances to be decomposed and some aspects of the initial experimental setup. Third, analogies guided Davy in the interpretation of the results of his experiments. When Davy's results were analogous to the results of an experiment that decomposed a particular substance, he infers that decomposition has once again taken place.

The roles that analogical reasoning played in Davy's work in electrochemistry were largely theory independent. One thing that is notably absent from Davy's discussion of analogy at the beginning of his textbook is any mention of theory. In the case of his work with the Voltaic pile, there was, at the time, no theoretical consensus regarding how the pile effected the chemical changes that it brought about. Despite the lack of theoretical consensus, Davy was able to use the pile to make various discoveries in part because he was guided by analogies. Hence, this case shows how analogical reasoning can guide scientific research in the absence of theoretical consensus.

By appreciating the roles of analogical reasoning in Davy's work, we arrive at a more complete picture of how analogical reasoning operates in scientific practice. Previous work on analogical reasoning has tended to focus on its role in modeling and theorizing. While this work is illuminating, Davy's analogical reasoning shows that we can't neglect roles of analogical reasoning that are less theory-focused and more directed towards experimentation.

Causation and Complex Phenomena: Econometric Modeling

Jennifer Jhun – Lake Forest College, United States

A careful investigation of history and practice reveals that econometric models are often not meant to be, strictly speaking, representational. Yet, they are expected to yield causal understanding by identifying the mechanisms underlying economic behavior. This may seem paradoxical; I argue that we can discharge these difficulties by paying attention to how econometricians

incorporate method into their models. These observations will have implications more generally for modeling complex phenomena, in particular more recent developments in multi-scale modeling, where analogies with engineering are salient.

Economics has become increasingly receptive to empirical work, where by empirical work I mean work that is informed by theory and/or data that partakes in the use of econometric or statistical tools. For instance, two prominent modelling strategies aimed at characterize the causal structure of a (macro-)economic system are structural and reduced-form modelling, the latter which belongs to a larger class of “non-structural” approaches. Structural models are systems of simultaneous linear equations that purportedly represent the underlying structure of the economy. But articulating them is difficult, requiring significant theoretical assumptions that, should they be mistaken, will most likely lack explanatory and predictive power. Reduced-form approaches instead involve systems of equations where all endogenous variables are functions of exogenous ones. This strategy is more manageable for experimental work (both “non-observational”, as well as natural and quasi-natural), doesn’t ostensibly require as many a priori assumptions, and is useful as first step in estimating time series data.

But while we can derive reduced form equations from an underlying structural model, we cannot go the other way. If we’re interested in generating reliable causal inferences in applied econometrics – which typically means policy analysis, i.e. counterfactual analysis – received wisdom tells us that what’s really needed is a structural model. That there is such a methodological distinction at all is an expression of a difficulty that econometricians today still struggle to bridge between the theoretical and empirical.

That is, we should not think, for instance, of non-structural approaches such as reduced or recursive modelling methods as inferior to structural ones by virtue of being incomplete or coarser grained versions of them. This historical triangulation will trace a line of thought that leads naturally to a live discussion in philosophy of science today (notably Batterman (2013) and Wilson (forthcoming)) concerning explanation when it comes to complex systems, in particular when such systems necessitate multiscale modeling. In particular, I claim that the econometrician face a problem that is analogous to what Batterman calls the “tyranny of scales,” which has actually been lurking in the history of econometric development over the past century. And in fact, it is familiar in one form to economists as the nefarious “problem of endogeneity.”

On the conditions for the objectivity of nutrition guidelines

Saana Jukola – Bielefeld University, Germany

This paper aims at assessing the conditions for producing objective nutrition guidelines. Objectivity is one of the main ideals of science. It is something we should aim at if we want to acquire trustworthy knowledge, i.e. knowledge that can be used for guiding actions in this complex world. But what is objectivity and what are its conditions? Recently philosophers have been active in examining and developing accounts of objectivity that can be used for assessing scientific practices (e.g., Douglas 2004; Daston & Galison 2007; Gelman & Hennig 2017). These accounts have demonstrated not only that objectivity is a complex concept but also that different understandings of what objectivity denotes can have practical implications as they guide our actions and influence methodological decisions. Consequently, it is important to be clear on what account one is committed to while evaluating scientific practices. In nutrition advice, the results of nutrition science are translated into recommendations, for instance Dietary Guidelines for Americans, which are meant to improve the health of the general public and to prevent chronic diseases. Recently the trustworthiness of nutrition advice has been questioned; established nutrition experts have been accused of having too close ties with the food industry. Apparent conflicts of interests are common in nutrition research, as big part of the field is funded by industry (e.g., Lesser et al. 2007). However, despite the identified problems with commercialized research, it would be much too simple to state that industry funding is a sign of research being biased. There are examples of sound, high-class privately funded research (e.g., Shapin 2008). Moreover, philosophers of science have argued against the value-free ideal of science: we should not assume that science should, or could, be free of so-called non-epistemic values (e.g., Longino 1990; Douglas 2009). The presence of these non-epistemic motivations can even be beneficial to science by creating more diverse research environments (Carrier 2010). But how to demarcate the acceptable influence of commercial interests from unacceptable if the goal is to produce trustworthy dietary guidelines?

I shall argue that in order to evaluate the trustworthiness of nutrition research and policy, we need an account of objectivity that takes into consideration the institutional conditions of research. By presenting examples from research on the relationship between sugar and health (e.g., Kearns et al 2015), I show that what is traditionally called the discovery side of science needs to be considered when knowledge production is evaluated and the conditions for objective guidelines assessed. This is because the way in which research projects and

questions are framed has a critical effect on what kind of knowledge is eventually available to inform decision-making. Thus, examinations of the conditions for objectivity should not focus solely on the conditions for justification procedures (for example, in the case of nutrition research, on how trials and observational studies are carried out) while disregarding how research can be skewed by extra-scientific factors.

Eliminating neuroscientific concepts of consciousness?

Yi-Hsuan Kao – National Yang-Ming University, Taiwan

Karen Yan – National Yang-Ming University, Taiwan

Irvine (2013) argues that the currently available neuroscientific concepts of consciousness should be eliminated based on some epistemological and pragmatic considerations. The aim of this paper is to argue that a scientific concept of consciousness is still needed even if this concept has some shortcomings that Irvine has already identified. In Section 1, I will summarize how Irvine argues that the currently available scientific concepts of consciousness need to be eliminated based on three pragmatic reasons. I will then present my argument against Irvine's argument in Section 2 and Section 3. The core of my argument is based on a fact that Irvine misses out an important pragmatic role played by a scientific concept of consciousness. In Section 2, I will use Feest's (2017) notion of object of research to elaborate the view that a scientific concept can play a pragmatic role in guiding researchers to describe and explore their object of research. Moreover, the boundary of the object of research is epistemically blurry in the sense that the relevant researches do not know the exact shape or contours of their object of research. Much of their empirical work is to delineate what they are actually interested in investigating (Feest, 2017). In Section 3, I will show that a scientific concept of consciousness can at least function as the concept of object of research as Feest has articulated. This gives a pragmatic reason for keeping a scientific concept of consciousness. In Section 4, I will further show that how it is possible for a scientific concept of consciousness construed as object of research to facilitate the integration of different scientific constructs relating to consciousness from different fields. This may give another pragmatic reason for keeping a scientific concept of consciousness.

Produce, Underlie, Maintain: What's Behind the Mechanistic Triad?

Lena Kästner – Ruhr-University Bochum, Germany

According to the mechanistic view, scientists explain phenomena by uncovering the mechanisms responsible for them (e.g. Machamer, Darden, and Craver 2000, Craver 2007, Bechtel and Abrahamsen 2005, Illari and Williamson 2012). What precisely this “being responsible for” means, however, is a matter of heated debate in contemporary philosophy of science (see e.g. Couch 2011, Harbecke 2010, Leuridan 2012, Glennan 2016). The suggestions vary significantly as an effect of what exactly we take to be “the phenomenon to be explained” and the “mechanism underlying it”, respectively. Most prominently perhaps, we are told to conceive of the relation between mechanisms and their phenomena in causal or constitutive terms depending on their spatio-temporal characteristics and boundaries (see also Kaiser and Krickel forthcoming). While such metaphysical analyses certainly shed light on phenomenon-mechanism relations, they remain silent about the relations between different explanatory projects in empirical research.

In practice, mechanisms are supposed to explain a whole collection of different things: processes with multiple phases, end products or final stages of processes, stable states, continuous or repetitive behaviors, even properties. Naturally the relationship between the “phenomenon” and the “mechanism responsible for it” will vary with the nature of the relata. A causal sequence can cause an end product to be generated, an arrangement of interactions can constitute an overall complex process, etc. While the explanatory projects focusing on causal and constitutive factors, respectively, are metaphysically quite different, scientists actually flip back and forth between them quite frequently even within a single research project. This is also highlighted by recent discussions on mechanism discovery. For instance, Craver and Darden (2013) suggest that researchers need to employ different strategies if they aim to find mechanisms that produce (cause), underlie (constitute), or maintain their phenomena, respectively.

While this may sound temptingly metaphysical, I suggest a somewhat pragmatic reading of the triad. Building on a well-known example from biology—lactose metabolism—I illustrate how scientists’ search for explanations is guided by different research questions throughout the discovery processes. Depending on exactly what research questions they ask at any given point, scientists will differentially emphasize causal, constitutive, or continuous aspects in the mechanistic explanations they construct. There are two lessons to learn from this: (i) Causal explanations are undoubtedly an important part of the enterprise of scientific explanation, but there is more.

Indeed, insights about constituents or maintaining factors may be more important for some explanatory projects. And, (ii) to speak of mechanisms that produce, underlie and maintain their phenomena, respectively, does not mean that there actually are three different sets of goings-on in the world. Rather, it is a matter of perspective. What exactly we take “the phenomenon” to be depends on the explanatory project we pursue and which research questions we ask. What kind of mechanism can be “responsible for it” is determined by the nature of the phenomenon. This, in turn, determines not only the phenomenon-mechanism relation at hand but also whether our mechanistic explanations will emphasize causal, constitute, or continuous aspects.

Bechtel, W. & Abrahamsen, A. (2005). Explanation: a mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 421–441.

Craver, C.F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. New York: Oxford University Press.

Craver, C.F. & Darden, L. (2013). *In Search of Mechanisms: Discoveries Across the Life Sciences*. Chicago: University of Chicago Press.

Couch, M. B. (2011). “Mechanisms and constitutive relevance”, *Synthese*, 83, 375–388.

Glennan, S. (2016). *The New Mechanical Philosophy*. New York: Oxford University Press.

Harbecke, J. (2010). Mechanistic Constitution in Neurobiological Explanations, *International Studies in the Philosophy of Science*, 24, 267–285.

Illari, P. & Williamson, J. (2012). What is a mechanism? Thinking about mechanisms across the sciences. *European Journal of Philosophy of Science*, 2, 119–135.

Kaiser, M. & Krickel, B. (forthcoming). The Metaphysics of Constitutive Mechanistic Phenomena. *British Journal for the Philosophy of Science*.

Leuridan, B. (2012). Three problems for the mutual manipulability account of constitutive relevance in mechanisms. *The British Journal for the Philosophy of Science*, 63, 399–427.

Machamer, P. K., Darden, L. & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1–25.

What are we doing when we describe something as being a part in a biological parts repository?

Catherine Kendig – Michigan State University, United States

Synthetic biology may be defined as the application of engineering principles to the design, construction, and analysis of biological systems. For example, biological functions such as metabolism may now be genetically reengineered to produce new chemical compounds. Designing, modifying, and manufacturing new biomolecular systems and metabolic pathways draws upon analogies from engineering such as standardized parts, circuits, oscillators, and digital logic gates. These engineering techniques and computational models are then used to understand, rewire, and reengineer biological networks. But is that all there is to synthetic biology? Is this descriptive catalogue of bricolage wholly explanatory of the discipline? Do these descriptions impact scientific metaphysics? If so, how might these parts descriptions inform us of what it is to be a biological kind? Attempting to answer these questions requires investigations into the nature of these biological parts as well as what role descriptions of parts play in the identification of them as the same sort of thing as another thing of the same kind.

Biological parts repositories serve as a common resource where synthetic biologists can go to obtain physical samples of DNA associated with descriptive data about those samples. Perhaps the best example of a biological parts repository is the iGEM Registry of Standard Biological Parts (igem.org). These parts have been classified into collections, some labeled with engineering terms (e.g. chassis, receiver) some labeled with biological terms (e.g., proteindomain, binding), and some labeled with vague generality (e.g., classic, direction). Descriptive catalogues appear to furnish part-specific knowledge and informational specificity that allow us to individuate them as parts. Repositories catalogue parts. It seems straightforward enough to understand what is contained within the repository in terms of the general concept: part. But what are we doing when we describe something as being a part? How do we know which part to use?, Which model do we follow?, In order to answer this set of questions, we need to be able to track parts—or at least the names of parts—as well as their diverse multilevel descriptions that are expressed in both formal labels and natural language descriptions.

I focus on the preliminary processes of knowledge production which are prerequisite to the construction or identification of ontologies of parts within synthetic biology. I investigate some problems arising from the varied descriptions of parts contained in different repositories. Following this, I

outline problems that arise with naming and tracking parts within and across repositories and explore how the comparison of parts across different databases might be facilitated. This focuses on computational models currently being sought that would allow practitioners to capture information and meta-information relevant to answering particular questions through the construction of similarity measures for different biological ontologies. I conclude by discussing the social and normative aspects of part-making and kind-making in synthetic biology and suggest that the activities associated with identification, description, and cataloguing in synthetic biology can be best understood as the storing of informational specificity that can be retrieved and used later.

Understanding, Accuracy, and the Aims of Science

Kareem Khalifa – Middlebury College, United States

Jared Millson – Agnes Scott College, United States

Accuracy monism is the idea that accurate representation (paradigmatically: the acquisition of true beliefs and the avoidance of false beliefs) is the only ultimate epistemic aim of scientific inquiry. Several authors argue that accuracy monism is false, and that understanding should either complement or supplant accurate representation as the ultimate aim of science. The arguments for this are threefold. First, past inquiries that resulted in false beliefs but advanced our understanding are episodes of scientific progress. Second, scientists' use of idealizations suggests that some falsehoods are cognitively valuable because they advance our understanding. Third, inquiries that aim at truths that do not advance our understanding appear deficient or misguided. This paper defends accuracy monism against these objections.

We argue that the first two objections fail to appreciate accuracy monism's insistence that accurate representation is science's ultimate epistemic aim. Such a view is compatible with falsehoods serving as more proximate aims and accruing instrumental epistemic value by serving as effective means for accurate representation. We argue that past theories are naturally interpreted along these lines by scientists. Using several examples from economics, we also argue that if idealizations do not accord with this accuracy-monist picture, then the understanding they are alleged to provide is of dubious epistemic value. We supplement these examples by showing how the leading critics of accuracy monism are tacitly committed to this picture.

We then turn to the last objection, which alleges that truths that fail to confer understanding are epistemically deficient. We argue that this objection ignores accuracy monism's insistence that accurate representation is science's

ultimate epistemic aim. Once this point is taken on board, the alleged counterexamples can be reinterpreted as failures to advance science's non-epistemic or pragmatic aims. We show how purveyors of this objection are committed to both an accuracy requirement and a pragmatic requirement, and that these two requirements suffice to account for what is deficient in the relevant cases. Hence, nothing further—such as placing understanding as one of science's ultimate epistemic aims—is needed to explain away the counterexamples.

This last defense suggests that far more of science's aims are pragmatic than accuracy monists typically acknowledge. Furthermore, scientists' interests in different questions provide a useful framework for exhibiting their different pragmatic aims. We use these points to highlight two interesting consequences of our view. First, while all scientific inquiries pursue questions that are correctly answered only with accurate representations, some of these answers do not advance scientific understanding. Second, inquiries that aim at understanding are generally guided by questions that seek information of high instrumental epistemic value and that serve several pragmatic interests. Combined, these two points explain away the intuition that understanding is an ultimate epistemic aim of science.

Synthetic Biology's Alternative Realities - Turning Fictional Systems into Concrete Ones

Tarja Knuuttila – University of Helsinki, Finland

Rami Koskinen – University of Helsinki, Finland

The notion of fiction has become a staple of recent philosophical discussion on modeling (e.g. Suárez 2008, Toon 2012, Levy 2015). Among other things, it has provided an answer to the ontological question of what kind of objects are nonconcrete models: if they are not to be identified with their targets or their material instantiations – e.g. particular tokens of mathematical equations on a blackboard – they have to be some kind of imaginary constructs that resemble works of literary fiction (e.g. Frigg and Nguyen 2017). Another reason to evoke fiction-talk is the realization that most models are, at best, highly unrealistic representations of reality, containing elements that are idealized, distorting or even strictly speaking nonexistent.

We distinguish two perspectives on fictional modeling that are often somewhat entangled in the current discussion: the case in which the model itself is fictional (i.e. an imagined object) and the case in which the target of modeling is a fictional (i.e. nonexistent) system. We concentrate on the latter case by studying two examples from synthetic biology. The practice of

synthetic biology brings to the discussion of fiction an interesting twist in that synthetic biologists also strive to turn the fictional alternatives to actual biological systems into concrete living entities and their parts. Such possible concretization of fictional systems shows, we argue, that it may often be more fruitful to approach fiction in terms of possibility than falsity.

As case studies, we focus on artificial biopolymers (Benner lab) and synthetic genetic circuits (Hasty lab). Synthetic biologists consider the study of such contingently nonexistent systems as an exploration of possible biology that goes beyond actual evolutionary designs. Could life be based on another molecule than DNA? Is it possible to rewire genetic circuits and metabolic networks of micro-organisms so that they produce new kinds of commercially valuable substances, even providing a basis for a biology-based industrial revolution? In these cases, the intended target system of a fictional model does not yet exist, apart as an unactualized possibility. However, such a fictional target may be rendered—often with a great effort—into an actual entity.

This role of fictions as tools for gearing the scientists' outlook towards the possible has not gained due attention from philosophers of science. For instance, the mechanists discuss about how-possibly models, but for them the ultimate research goal consists of delivering a how-actually explanation. In mechanistic discussion, how-possibly explanations are portrayed in terms of their lack of detail, rendering them mechanism schemas at best. In our view, models lacking details should not be conflated with models aimed at studying possibilities. Philosophy of science has invested most of its efforts in analyzing how scientists successfully represent and model of what there is. In contrast, synthetic biology furnishes an example of a scientific field, whose modeling activities are decidedly geared towards what there could be. The question to be asked is whether or not many other modeling practices trading with fiction, even outside of the larger expanse of engineering sciences, are also of this character.

Interpreting Archaeological Material: The Model of Evidential Reasoning Extended

Kristin Kokkov – University of Tartu, Estonia

Archaeology is a domain that studies material remains of past events for the purpose of understanding social structures and cultural dynamics of past people. The events and people in question do not exist anymore and cannot be observed directly. Thus, there is a gap between the subjects that are studied and the information that is preserved from the past.

Archaeologists have tried to overcome this interpretational gap for many decades. In the 1970s, Lewis R. Binford introduced the method of middle-range theories as a tool of archaeological interpretation. In the 1980s, Ian Hodder laid the foundation for the post-processual movement that emphasised the importance of understanding past social context in interpreting past material culture.

In recent years, the question of the interpretational gap between material remains and past events has been analysed by Alison Wylie. To explain how archaeologists interpret material remains, she (2011: 371) suggests the model of evidential reasoning. Wylie (2011: 380) describes this model by saying that it involves three functional components: 1) empirical input; 2) theory that mediates the interpretation of empirical input as evidence; and 3) the claims on which this empirical input bears as evidence.

Taking this model as the basis for my study, I examined the archaeological research process in detail. It became evident that the process of archaeological interpretation is somewhat more diverse than the model of evidential reasoning *prima facie* suggests.

Therefore, I propose an extended version of the model. I claim that the process of archaeological theory formation consists of at least three different stages of interpretation that proceed from the present material remains towards the past events:

- 1) the stage between material remains and archaeological record;
- 2) the stage between the description of the archaeological record and claims about the past;
- 3) the stage between claims about the past and general theory about the past historical-cultural context.

I argue that each of these stages has the structure of the model of evidential reasoning, but has its own specific function. In the first stage, the material remains are interpreted as archaeological record. In the second stage, archaeologists make claims about the past and explain why the archaeological record is the way it appears to us. In the third stage, archaeologists try to explain why these past events took place that left behind the archaeological record we can see today.

My aim is to explain in detail the structure of each interpretational stage, and show how the archaeological research gradually proceeds from the material remains towards the understanding of the cultural past.

Wylie, Alison 2011. *Critical Distance: Stabilising Evidential Claims in Archaeology*, in Philip Dawid, William Twining, Mimi Vasilaki (Eds.) *Evidence, Inference and Enquiry*. Oxford, New York: Oxford University Press, 371–394.

With a little help from my (old) friends? Evaluation of evolutionary explanations of diseases

Nina Kranke – University of Münster, Germany

Within the last couple of decades, evolutionary explanations of diseases have become increasingly significant. An example for evolutionary explanations in medicine is the so-called old-friends hypothesis that explains increasing incidences of autoimmune diseases in industrialized countries. The hypothesis suggests that the lack of exposure to various organisms (e.g. helminths) that have accompanied mammalian evolution for a long period of time and had to be tolerated, could have detrimental effects. In this paper, I use the discussion of the old-friends hypothesis as a case study to analyze how different interest groups evaluate evolutionary explanations of diseases. Since the discussion extends beyond the realm of medical research, I also include other interest groups such as science scholars, science journalists, companies, and patients in my investigation.

My analysis shows that different groups come to different conclusions about the value of the old-friends hypothesis, because they have different interests and apply different evaluation criteria. While some science scholars argue that ultimate explanations of diseases are speculative and irrelevant, many medical researchers seem to regard them as plausible, at least in combination with proximate explanations. Science journalists and patients also know of the old-friends hypothesis and most of them consider it plausible. In some cases, their confidence in the hypothesis and/or their desperation goes so far that they infect themselves with worms to treat their autoimmune diseases. This active decision to host worm colonies in order to reshape the immune system creates a sense of autonomy and independence from the health system. Since, in the US, it is illegal to sell worms for therapy, the New York Times Magazine calls this phenomenon a “shadow network of patients” and the New Scientist writes that “scientists need to catch up” and claims that “it is surely time to loosen the regulations and encourage fruitful collaborations between scientist and citizens”.

By focusing on the epistemic value of individual evolutionary explanations in medicine, science scholars have overlooked their social function. I argue, that evolutionary explanations can play an important role in intra and interdisciplinary integration of scientific results. This function is particularly relevant for studies of host-parasite interactions which are situated at the intersection of biology and medicine with a lot of potential for explanatory integration. The relatively long history of evolutionary explanations in fields that are interested in host-parasite systems is another reason why researchers

in these fields are inclined to take this type of explanation seriously. My analysis also raises questions about science communication and social responsibility of scientists. In the information age, patients have easy access to scientific articles which are considered trustworthy sources. Since research on immunomodulating abilities of helminths is still in its infancy and the old-friends hypothesis has not yet been sufficiently confirmed, it is important that scientist and science journalists communicate scientific results and hypotheses accordingly. The case study shows once again that the boundaries between scientific communities and other groups are fuzzy and invites policy-makers, science scholars, and scientists to think about possibilities of integrating knowledge produced by scientists with knowledge produced outside of scientific contexts.

Randomised controlled trials: The biases and limits arising in practice

Alexander Krauss – London School of Economics, United Kingdom

Evidence from randomised controlled trials (RCTs) is generally thought of as the most reliable form of evidence within the medical and social sciences. The RCT method is viewed by most scientists and many medical practitioners and policymakers as having revolutionised the medical sciences in the second half of the 20th century and many of the behavioural and social sciences since then. Researchers in these fields typically claim that randomised experimentation is the best and often only means to generate rigorous ‘causal’ knowledge. They also often think of RCTs as being free from overly complex theory and strong methodological assumptions and biases that unavoidably affect other methods.

To better understand this leading scientific method and its limits, philosophers of science have been increasingly studying the theory behind RCTs. But philosophers largely use abstract reasoning to discuss issues related to randomisation, statistical probabilities, deductive reasoning, ethical implications and the like. Worrall (2007; 2007a) for example assesses the function and limitation of randomisation and outlines some of its ethical constraints. Cartwright (2007; 2010) reasons about statistical probabilities, external validity and the conditions under which causal conclusions follow deductively in the ideal RCT. Clarke, Gillies, Illari, Russo and Williamson (2014) argue for a rethinking about hierarchies of evidence and the position of RCTs within them, and for a greater focus on mechanisms and not just correlations as evidence for causal claims. These leading philosophers, while making important contributions in improving our understanding about the RCT

method, largely take a theoretical perspective and do not systematically study the implementation of RCTs in the real world.

This paper instead takes a broader and more applied approach by combining philosophical, methodological and scientific perspectives and by providing concrete examples from the ten most cited RCT studies that together can significantly improve our understanding of the RCT method. The paper thereby outlines a large number of new and important theoretical assumptions, methodological biases and empirical limitations not yet discussed in the medical, social or philosophical literature that emerge when designing, implementing and analysing trials in practice. These assumptions include that participants' background traits that affect outcomes would not change between trial groups during trial implementation (i.e. not just no differences at baseline but also at endline); that the particular time points for the baseline and endline would be chosen to adequately reflect the average (or greatest possible) treatment outcome; that randomisation can ensure participants are evenly distributed between trial groups along measurable, non-measurable and unknown background influencers, among many others.

This paper thereby provides a much more comprehensive overview of the range of issues and problems facing RCTs (which is needed to assess an RCT's overall robustness) than in the existing literature that tends to focus on specific issues. Epistemologically, the paper shows that RCTs generally have some degree of bias in their results – as illustrated by assessing the ten most cited RCT studies worldwide. A central and important epistemic topic underlying the paper is thus whether or not RCTs can, despite the range of issues and problems associated with them, provide sufficiently rigorous evidence that would allow us to be confident in their reported causal claims. The paper's implications include that we must not overly rely on any single research method but always combine methods.

What's Wrong with (the Recent) Criticism of Research on Gender and Racial Biases?

Anna Leuschner – Leibniz Universität Hannover, Germany

The paper takes issue with recent criticism of research on gender and racial biases. I examine Allen-Hermanson's criticism of the "shooter bias" (a potential explanation for the disproportionate number of minorities killed by the U.S. police force), and criticism put forward by Ceci, Williams, and colleagues regarding the situation of women academics.

First, the points raised by Allen-Hermanson (2017), and Ceci, Williams, et al. (2014; 2015) are methodologically problematic. Allen-Hermanson claims that

the shooter bias is not empirically confirmed while there is substantial counter-evidence. However, I will show that this does not stand up to close scrutiny. Ceci et al. claim that there is no gender bias in academia anymore, because (a) they did not find any gender bias in manuscript and grant proposal assessments and (b) at the top academic level women job candidates are preferred over men. I will discuss problems of both points.

Second, these points of criticism are presented polemically. For example, Allen-Hermanson unwarrantedly accuses authors in Brownstein and Saul's compendium "Implicit Bias and Philosophy" of being biased and focusing on "basically irrelevant" issues; and Williams goes on record as saying: "It's tempting to blame gender when you don't get a job and you're a woman [...]. It's easier ... than to admit that the entire premise of what you've done for the past 7 years of your life was flawed at the root." (quoted after Benderly 2015)

As is known from other contexts, academic debates can be seriously affected by hostile criticism. The probably best examined field in this respect is climate science, where scientists have been targeted by personal attacks from climate change deniers: they have been confronted with the never-ending manufacture of pseudo-evidence claimed to disprove their findings, and they have experienced continual doubts with regard to their competence and reliability.

By drawing on Biddle et al. (2017) I will argue that hostile polemics can have serious epistemic effects: it can stifle discussions as intimidated researchers become reluctant to address certain questions or hypotheses, and it can lead to a lopsided distribution of inductive risk. Thus, while criticism is actually key for epistemic progress using such polemical style in criticizing research is epistemically problematic.

Allen-Hermanson, Sean (2017). Book Review of "Implicit Bias and Philosophy. Michael Brownstein and Jennifer Saul (eds.), OUP 2016." *Philosophy* 92 (2), 315–322.

Benderly, Beryl Lieff (2015). Women Have a Hiring Advantage in the Scientific Stratosphere. *Science Online*:

<http://www.sciencemag.org/careers/2015/04/women-have-hiring-advantage-scientific-stratosphere>.

Biddle, Justin, Ian James Kidd, and Anna Leuschner (2017). Epistemic Corruption and Manufactured Doubt: The Case of Climate Science. *Public Affairs Quarterly* 31 (3), 165–187.

Ceci, Stephen J., Donna K. Ginther, Shulamit Kahn, and Wendy M. Williams (2014). Women in Academic Science: A Changing Landscape. *Psychological Science in the Public Interest* 15 (3), 75–141.

Williams, Wendy M. and Stephen J. Ceci (2015). National Hiring Experiments Reveal 2:1 Faculty Preference for Women on STEM Tenure Track. *PNAS* 112 (17), 5360–5365.

Inference, Norms, and Model-Based Understanding

Harry Lewendon-Evans – Durham University, United Kingdom

In recent decades, philosophy of science has paid significant attention to the use of models in science. This focus on models and the activities involved in modelling have been broadly driven by an interest in scientific practice, as opposed to formal reconstructions of scientific knowledge (Morgan and Morrison 1999; Weisberg 2007; Godfrey-Smith 2006). A more recent development has argued that the epistemic value of models lies in the ability to provide understanding (de Regt, Leonelli and Eigner 2009; Gelfert 2016). However, although it is generally agreed that understanding is a significant epistemic virtue of modelling practice, there is yet no agreed characterization of model-based understanding or which aspects of modelling practice are necessary for understanding (Psillos and Nounou 2012).

This paper addresses this issue by articulating and defending an inferentialist approach to model-based understanding. The inferentialist approach holds that models function as external, inferential aids for facilitating understanding, as recently proposed by Kuorikoski and Ylikoski (2015). However, drawing on a range of concrete examples from the history and practice of science, this paper extends the inferentialist account of model-based understanding by identifying three ways in which models facilitate inferential understanding which have so far not been accounted for.

First, I argue that failed models provide understanding insofar as they help to articulate material inferential incompatibilities. These refer to the class of inferences that are prohibited given the shared, theoretical commitments of a research community. Second, I highlight the counterfactual or modal nature of modelling practice, arguing that models facilitate understanding by delimiting the range of inferences that are possible given these shared, theoretical commitments as well as providing the basis for the generation of how-possibly explanations (Bokulich 2014). Third, I argue that there is an important normative dimension to model-based understanding that has yet to be acknowledged in the literature. There are two aspects to this normative dimension corresponding to distinct stages of the modelling process: (a) in the construction of a model, where we see the establishment of inferential rules, which legitimise inferential moves from claims about the source to claims about the target; and (b) in the acceptance of a model within a research

community, where models function to sculpt particular patterns of reasoning and subsequently ‘perpetuate communal norms of intelligibility’ (Woody 2015: 81). In this respect, I argue that attending to the normative function of modelling in scientific practice provides a novel way of conceiving model-based understanding.

In conclusion, this paper sheds light on the characterisation of model-based understanding by identifying three additional ways in which models underpin our inferential reasoning: (i) material incompatibilities; (ii) counterfactual inferences; and (iii) the generation of inferential rules and how these inform a research community’s norms of intelligibility. This paper therefore significantly extends an inferentialist approach to model-based understanding by providing a finer-grained characterization of the kinds of inferences involved. By introducing the connection between scientific models and shared norms of intelligibility of research practices, this paper proposes a new way of analysing the function of models in scientific practice.

On the Evolutionary Synthesis of the Social Sciences: A Philosophy of Social Science in Practice Perspective

Simon Lohse – Leibniz Universitaet Hannover, Germany

In recent years there has been a renewed interest in attempts to synthesize (or unify) the social sciences using evolutionary thinking. In this talk, I want to discuss two central questions in this context: (1) Is a theory of cultural evolution a good candidate to synthesize the social sciences? (2) What is the added value of evolutionary explanations for the social sciences? My aim is to highlight some hitherto underestimated challenges for evolutionary approaches that come into view when one looks at these questions against the backdrop of actual scientific practice in sociology and political science, arguably two centrepiece disciplines of the social sciences. However, instead of rejecting an evolutionary synthesis of the social sciences on principle grounds or to argue for the lack of explanatory power of Darwinian thinking in the social sciences (Schatzki 2001), I will make the positive case for more interdisciplinary dialogue that is sensitive to the epistemic particularities of the social sciences. The first part of my talk will scrutinize five background assumptions of one of the most promising evolutionary candidates for synthesizing the social sciences (Mesoudi et al. 2006; Mesoudi 2011):

- (1) Pluralism is bad for scientific progress.
- (2) All of the social sciences investigate the same “cultural stuff” at the end of the day.

(3) The (poor) state of the art in the social sciences is mainly due to the lack of an integrative theoretical framework.

(4) Social scientists reject formalized evolutionary models because they do not like idealizing/simplifying models.

(5) Evolutionary biology and the social sciences share an epistemic core interest.

I will show that, most notably, assumptions (2), (3) and (5) are problematic due to ontological incommensurability, the lack of a corroborated explanation of the multiparadigmatic state of the social sciences, and the neglect of the wide variety of epistemic interests and practices of sociologists and political scientists. I will argue that the concurrence of these challenges threatens the success of a synthesising evolutionary approach to the social sciences.

The second part of my talk will dovetail with my analysis of assumption (5) and address the question about the added value of evolutionary approaches to the social sciences. I will discuss a number of recent empirical studies from top social science journals to argue for a level-headed answer to this question: Evolutionary models of cultural phenomena are not better or worse tout court but they are one useful epistemic tool amongst others. I will conclude my talk with a few recommendations for increasing the acceptance of evolutionary approaches in the social sciences.

Mesoudi, A. (2011). *Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences*. Chicago; London: University of Chicago Press.

Mesoudi, A., Whiten, A., & Laland, K. N. (2006). Towards a Unified Science of Cultural Evolution. *Behavioral and Brain Sciences*, 29(4), 329–347.

Schatzki, T. R. (2001). On Sociocultural Evolution by Social Selection. *Journal for the Theory of Social Behaviour*, 31(4), 341–364.

‘Constituting’ tension? An epistemological analysis of the role of measurement and coordination in Ohm’s scientific practice

Michele Luchetti – Central European University, Hungary

In this paper, I assess and develop the current views on the ‘problem of coordination’ in relationship with measuring practices. The problem of coordination has been recently characterised as the issue of how measurement procedures can justifiably be said to measure the parameter of interest in the absence of independent ways of assessing them (Tal, 2013). Two major accounts characterise the process of coordination in terms of mutual

refinement of theory and measurement standards (Chang, 2004; Van Fraassen, 2008). Although they place a different emphasis on the extent to which measurement procedures can develop independently from theories, both Chang and Van Fraassen claim that quantity terms are 'constituted' through the historical process of mutual coordination.

The achievement of a successful coordination involves various actors (instruments, models of measurement, theoretical and metaphysical assumptions, etc.) at different stages of development, and understanding the fixing of constitutive components can involve several levels of analysis. These can be abstracted away, to distinguish only between a top-down constitutive role of certain theoretical principles, and a bottom-up constitutive role of measurement procedures (Padovani 2017), but can be further specified, depending on the details of specific scientific inquiries. I claim that, in the attempt to clarify what gets constituted, and how, along the process by which a coordination is achieved, degrees and dimensions of 'constitutivity' can emerge. Thus, there is a different sense by which a measurement outcome, a quantity term, and an empirical law are constituted during this process.

I support this claim by analysing Ohm's scientific work on electric conductivity. Although in Ohm's early experimental papers there is no indication about how to precisely measure neither tension nor electroscopic force – both figuring as parameters in his formula – the use of the thermocouple in his measurement apparatus allowed him to test for actual electroscopic force, rather than loss of force. Justification for this came from the assumption that difference in electroscopic force is linearly correlated with temperature differential, which is what was kept constant by the thermocouple in the measurement apparatus. Given the epistemic role played by the measurement apparatus and by the assumption of linearity, the measurement outcomes of electroscopic force can thus be considered as 'constituted'.

In addition, in Ohm's (1826) mathematical formulation of the relationship between resistance and current, he identified tension and current electricity, which to his contemporaries referred to two distinct types of phenomena (static the one, current the other). To make sense of his experimental results, he posited a tension between two adjacent elements of a closed circuit, while the received view considered it as a possible property of an open circuit only. Thus, the quantity term 'tension' was constituted by both the procedure used to measure electroscopic force, and a theoretical assumption. Such assumption was finally incorporated by Ohm in the notion of tension (*Spannung*) in his 1827 mathematical treatise, where it was 'elevated' to the status of a fundamental principle (Schagrin 1963; Archibald 1988): only at this stage, it can be said that the empirical law was established.

More Talk About Toy Models

Joshua Luczak – Leibniz Universität Hannover, Germany

Toy models are models that are not intended to perform a representational function, but rather to perform some other important function. For example, scientists often use them:

1. To learn to use, or to become comfortable with, certain formal techniques (e.g. renormalization). That is, as a pedagogical device.
2. To elucidate certain ideas relevant to a theory. That is, to reach a clearer understanding of an idea, its implications, and its relation to other ideas within a theory.
3. To test the compatibility of various concepts (i.e. in a consistency proof).
4. To generate hypotheses about other systems.

One commonly finds authors using simple models to perform one or more of these functions in the introductory chapters of physics textbooks. When they do so, and do not intend for them to perform a representational function, it is appropriate to regard their models as toy models. Some examples common to statistical mechanics include: Tatiana and Paul Ehrenfests' urn and wind-tree models, Mark Kac's ring model, the baker's transformation, the Ising model, and the Arnold cat map.

Philosophers are mostly concerned with representational models. In fact, most of the literature on modelling in science is concerned with categorising representational models into distinct types, with determining their ontological status, with explaining how it is that they achieve their representational function, and with articulating their relationship to concepts such as explanation, understanding, simulation, and approximate truth.

Toy models, in contrast, have received little attention from philosophers. This is despite their importance for science, their distinct nature, their frequent use, and the fact that they raise important and unique philosophical questions. Questions, such as: how can they be used to learn or generate hypotheses about features of the world when they do not represent real world systems?

A notable exception is the recent work of Joshua Luczak. Luczak drew attention to the distinct nature and importance of toy models. He did this by distinguishing them from approximations and idealisations, by highlighting and elaborating on several ways the Kac ring is used as a toy model, and by explaining why toy models can be used to successfully carry out functions 2-4, listed above, without performing a representational function. Luczak also encouraged a much greater philosophical discussion of toy models, so as to further our understanding of them, their role in science, and their relation to other model-types.

This article aims, among other things, to continue this discussion by elaborating on and furthering several parts of Luczak's work, and by discussing another model that is typically not intended to perform a representational function, but is rather used to perform some other important functions: Tatiana and Paul Ehrenfest's urn model. While the model was originally introduced so as to reason about the kinetic theory of gases and Ludwig Boltzmann's original attempts to account for irreversible thermal phenomena and the Second Law of Thermodynamics, the model has since been used to contribute to the fields of biology, chemistry, computer science, ecology, economics, and psychology. This article intends to showcase ways in which the urn model can be used to successfully perform functions 1-5, and it intends to justify why it can be used in these ways without performing a representational function. More specifically, this article intends to add to the discussion (i) by highlighting that toy models, so understood, do not perform a representational function on any of the leading substantive accounts of scientific representation, (ii) by offering reason to think that an agent's intention to represent is necessary for a model to perform a representational function, (iii) by highlighting and explaining how the urn model can and has been used to construct a more sophisticated model (Mark Kac's ring model) without intending that it perform a representational function, and (iv) by highlighting and explaining how the model can be used to learn to use, or to become comfortable with, a certain formal technique: the Method of Steepest Descent.

Mechanisms and Holism in Traditional Ecological Knowledge

David Ludwig – Wageningen University, Netherlands

Luana Poliseli – Universidade Federal da Bahia, Brazil

Conservation biologists and ecologists increasingly recognize the epistemic and political importance of engaging with Traditional Ecological Knowledge (TEK) of local communities. While there has been a quickly growing literature on the integration of TEK and Western science, it is also becoming increasingly recognized that the ideal of integration does often not translate into successful and harmonious collaboration. Causes of integration failures are complex and tend to involve political factors such as inequity in environmental governance and policy as well as epistemic challenges such as reliance on different methodologies and sources of evidence.

Debates about the epistemic limitations of knowledge integration commonly build on a contrast between the “holistic” character of TEK and the “mechanistic” orientation of Western science. The aim of this article is to critically engage with this contrast through the current literature on

mechanisms in philosophy of science and two case studies of TEK in practice. First, we introduce the case of water management and pest control in Balinese rice farming. Balinese TEK illustrates that local communities often have resources for identifying mechanisms of complex ecological dynamics and for intervening in these mechanisms through adaptive management strategies. A simple dichotomy between the holism of TEK and mechanistic character of Western science is therefore inadequate and runs the risk obscuring epistemic resources that are provided by TEK as well as marginalizing holders of TEK in conservation practices.

While holders of TEK often identify ecological mechanisms, holism debates also convey important insights about the limitations of mechanistic approaches in TEK. Our second case study focuses on the monitoring and management of caribou populations among Chisasibi Cree in the North American subarctic. While Cree have developed sophisticated strategies for monitoring and managing caribou populations, they do not involve mechanistic analyses that decompose target phenomena into interacting parts and other entities. Instead, Cree TEK is successful in part because it employs indicators and rules that sidestep questions about precise mechanisms of complex ecological dynamics in favor of flexibility and applicability in hunting and other community practices. While this case study indicates limitations of mechanistic approaches in TEK, we argue that similar trade-offs between mechanistic precision and applied significance are common in Western conservation management and applied ecology.

Taken together, the two case studies suggest a nuanced picture that challenges a simple distinction between holistic and mechanistic methodologies in TEK. Holders of TEK are perfectly capable of identifying ecological mechanisms but it is also true that TEK often deals with complexity by sidestepping issues of mechanisms in favor of more flexible indicators and rules. While the situation is not that different in Western ecology, there remain differences that become especially relevant if “holism” is understood in a wider sense that includes epistemology but also more general cultural aspects such as Indigenous emphasis of interconnectedness. We conclude by arguing that the proposed analysis can contribute to a framework that takes the epistemic resources of TEK seriously without downplaying substantial differences between Traditional and Western perspectives on ecology.

The Personal Equation in Astronomy: Triumph of Psychology or the Progress of a Fudge Factor?

Matthew Lund – Rowan University, United States

In 1796, Astronomer Royal Nevil Maskelyne fired his assistant, David Kinnebrook, for “observing the times of the Transits too late.” In 1823, after an analysis of the Royal Observatory’s data, F.W. Bessel published a startling conclusion: different observers of astronomical phenomena detect transits with ‘involuntary constant differences’ in the times of the observed events.

Bessel’s thesis eventually led to the idea of the ‘personal equation’ for observers, and spurred psychological investigations into the processes of visual perception. According to the dominant narrative (Boring 1929), empirical psychology’s development of the personal equation put observational astronomy back on an objective footing. However, this paper argues that practical innovations within observational astronomy itself led to the stabilization of data, and that robust psychological accounts of involuntary perceptual differences only emerged after this had occurred.

This paper investigates the reasoning process Bessel went through in his discovery of constant differences and asks whether such apparent perceptual relativity was viewed as a threat to objectivity. The contemporary reactions to Bessel’s report were tepid and sparse: “A few footnotes, a new column added to the tables of observations and two or three sentences are all that testify to the initial reaction to Bessel’s findings.” (2007, 354). Nonetheless, Bessel’s discovery revealed that the epistemic terrain of astronomy was much more unpredictable than had been previously thought. Yet the solution to these worries was not a rigorous theory of the observer, but rather a set of cautionary practices in data recording. As Hoffman characterized it, the discovery of constant differences brought into being a ‘cold tradition’, wherein epistemic problems are “preserved in the form of undiscussed practices.” (356)

Astronomical data and practice were precious commodities. Bessel was not concerned to provide a complete picture of observational psychology. He only wanted to supplement astronomy with a minimal epistemic account of observation so that the historical practices of astronomy could be preserved and extended. The ‘observer as instrument’ silently enters the picture with Bessel, but the observations are the item of interest, not the observer. In general, this paper argues that forms of scientific practice can act as structures of epistemic support, even in advance of a tenable (and conscious) epistemology of observation.

- Bessel, Friedrich Wilhelm. 1823. *Astronomische Beobachtungen zur der Königlichen Universitäts-Sternwarte in Königsberg von F.W. Bessel*, 8. Abtheilung vom 1. Januar bis 31. December 1822, Königsberg, III—VIII.
- Boring, E. G. 1929. *A history of experimental psychology*. New York: Appleton-Century.
- Canales, Jimena. 2009. *A tenth of a second: a history*. Chicago: University of Chicago Press.
- Chapman, Allan. 1983. The accuracy of angular measuring instruments used in astronomy between 1500 and 1850. *Journal for the History of Astronomy* 14 (2): 133-137.
- Hoffmann, Christoph. 2007. Constant differences: Friedrich Wilhelm Bessel, the concept of the observer in early nineteenth-century practical astronomy and the history of the personal equation. *British Journal for the History of Science* 40(3): 333-365.
- Maskelyne, Nevil. 1799. *Astronomical observations made at the Royal Observatory at Greenwich from 1787 to 1798 by the Rev. Nevil Maskelyne*. London.
- Mollon, J.D. and A.J. Perkins. 1996. Errors of judgement at Greenwich in 1796. *Nature* 380: 101-102.
- Sanford, E.C. 1888. Personal equation. *The American Journal of Psychology* 2(1): 3-38.
- Schaffer, Simon. 1988. Astronomers mark time: discipline and the personal equation. *Science in Context* 2: 115-145.

Economics imperialism as an instance of scientific imperialism: epistemic advancement, abuse of power, or both?

Magdalena Matecka – University of Helsinki, Finland

The aim of my paper is to contribute to the debate on scientific imperialism (SI) and on economics imperialism. First, I propose my account of SI. Afterwards I show how it can advance analysis of economics imperialism and why it matters.

The philosophy of science debate on scientific imperialism has revolved around the question of the permissibility of the application of scientific theories and methods outside the discipline in which they were initially introduced. Dupré characterizes SI as an application of a “successful scientific idea” “far beyond its original domain” (Dupré 2001, p. 74), so that this application cannot “provide much illumination”. For Clarke & Walsh (2013) SI is illegitimate occupation by one discipline of another discipline’s territory. My proposal

builds upon Mäki's (2013) notion of imperialism of standing, as well as it accommodates Dupré's and Clarke&Walsh's intuitions that there is something normatively problematic about SI.

I argue that some novel application X of methods, theories, research programs becomes imperialistic when:

1) X is favoured (by members of the scientific community) at the expense of other methods, or theories, or research programs in terms of academic and non-academic prestige, power, resources;

2) by claiming that X is...

a) more 'progressive' than applications of other methods, or theories, or research programs,

b) more 'scientific' than applications of other methods, or theories, or research programs;

3) and claim (2) is assumed to hold without providing argument for it.

Thus, in my view, SI is an activity that is related both to a certain view on progress (the epistemic aspect), as well as to a power to realize it -it is in fact favoured (the institutional aspect). This power manifests itself in the ability to affect standing between scientific approaches, at the same time without providing an argument why the approach of 'imperializers' is epistemically more advanced (progressive). This ability is conditioned by the power position of the institutional discipline from which the imperializing research approach originates.

I bring my account of SI to the analysis of economics imperialism. It seems that in the debate on economics imperialism the different sides have been talking past each other each other. The proponents of economics expansion, the so called 'economic imperializers', have been emphasizing only the epistemic aspects of applying economics outside its domain and claiming it is progressive, for example based on the idea of integration. The critics, have been mostly pointing out to the power issues and struggles, without engaging into epistemic discussion as such. Both aspects should be evaluated, and separated. Maybe even the possibility of both arguments being correct should be taken seriously: it is possible that including an approach from economics brings important epistemic progress while it taking over institutional resources is diminishing the possibility of some other approaches to develop into their full epistemic potential.

I will illustrate my point by analysing two case studies of economics imperialism recently discussed in the literature on SI – Becker's economics of discrimination (Chassonnery-Zaigouche 2018) and economic analysis of oppression (Rolin 2018).

Paraconsistent Heuristics, Inconsistent Information and Scientific Practice

Maria Del Rosario Martinez Ordaz – UNAM, Mexico

It is commonly believed that the study of contradictory beliefs, and the different ways in which epistemic agents manage and resolve such contradictions, plays a fundamental role in elucidating the foundations of rationality (Rovane 2004). For that reason, during the last four decades, philosophers and logicians of science have paid special attention to the different ways in which scientists can handle inconsistent information (Smith 1988; da Costa 2000; Batens 2002; Meheus 2002; Priest 2002; Brown and Priest 2004, 2015; Brown 2016a, 2017). The growing interest in the inconsistent character of scientific reasoning has given rise to, at least, two different types of research projects: the paraconsistent logics approach, and the paraconsistent reasoning strategies approach (Batens 2000; Brown and Priest 2004, 2015; Brown 2016a, 2016b, 2017). On the one hand, the former is mostly focused on the analysis of different types of logical consequence, and because of this, it has been claimed that it systematically overlooks the actual phenomenon of handling inconsistency in scientific practice. On the other hand, the latter is especially interested in analyzing general procedures that help to attain reliable information through the use of inconsistent information, and even though paraconsistent reasoning strategies often substantiate the general dynamics of certain logics, they are -most of the time- also logic-independent.

Here, I present a way to understand the practices of handling inconsistent information in science as the use of paraconsistent reasoning strategies as heuristics. In order to do so, I proceed as follows. First, I introduce some of the most important critiques to the paraconsistent logics approach, I pay special attention to the objection of how the logical reconstructions provided by this approach are not explicative of the actual practice of handling inconsistency in science (Vickers 2013, Boccardi and Macias-Bustos 2017). Second, I present the paraconsistent reasoning strategies that are already available in the literature and explain their differences –mainly regarding the type of inferential procedures that they emphasize and the inferential contexts in which they are optimally used. Third, I claim that such reasoning strategies behave as heuristics in the sense of being regularities for actions wherein “a kind of action (behavior) is characteristically undertaken under specifiable kinds of circumstances to achieve an end, or as part of a larger plan that is designed to do so” (Wimsatt 2007, p. 346). Finally, I argue that the use of these strategies,

as heuristics, could be explicative of the ways in which scientists often handle inconsistency in their practices.

Two Dogmas of Representationalism

Dana Matthiessen – University of Pittsburgh, United States

In recent years a standard view of scientific representation has been promoted by figures such as Teller, Godfrey-Smith, Pincock, Giere, Frigg, and structural realists. While each view has its individual differences, all subscribe to the notion that (i) the representational capacity of a successful scientific theory consists in a universally-instantiated relation of isomorphism, homomorphism, or similarity between elements of a theoretical model and features of the world, and (ii) that scientists test the accuracy of a given model by directly comparing it to the features of the world. Stated at its usual level of generality, I find this account unsatisfactory. This notion of representation explains the success of theories in terms that presuppose that features of the world are arranged in a manner that allows for a simple comparison between the world and a model. This is circular if we only access such features of the world with the help of these same models. But how else could such a comparison be made? Answering this question calls for a closer look at how models are related to scientific experimentation.

Using examples from protein science and particle physics, I will discuss the ways that theories are employed in relation to experimental practice. The first way involves the manner in which experimentalists draw on theories as locally applicable pools of information (Cf. Waters)—a patchwork of partial theories, technical know-how, and practical considerations that together provide the core concepts, strategies, and concerns by means of which researchers make the specific decisions and perform the particular actions that take them from the initial preparation of a system to a final result. It is by means of this local theory that they understand what they are doing in each experimental niche and what they cannot do there.

This body of knowledge is drawn on not only to execute an experiment, but also in the processing of the resulting data. Here the initial results are subjected to various techniques of data-shaping meant to mold data expressing the highly idiosyncratic conditions of a particular experiment into a more general form. It is through this latter form, shaped by an understanding of local conditions, that experimental results can be interpreted and “compared to” broadly applicable theoretical models. But this is a complex process in which the data model may be altered according to other bits of theory with empirical backing. There is no getting “outside” of the models to

the world here; while this process may point to significant limits in the inferential utility of a given model, it does not permit of a direct model-world comparison capable of yielding definitive confirmation.

From this point of view, questions about the representational capacity of theories tend to boil down to the problem of reliably coordinating a great range of technical practices. Theories enable this through the two roles discussed above: by informing local data gathering and processing practices, and by enabling inter-model comparisons that allow diversely obtained data sets to “talk to” one another.

Scaffolding the Science Behind Forensic Science

Barton Moffatt – Mississippi State University, United States

The forensics science community in the United States has begun to grapple with an ongoing series of faulty forensic science scandals. The Federal Bureau of Investigation recently revealed that 20 years of testimony matching hair samples was flawed. Fire science has undergone a radical transformation upending years of received wisdom. Other sub-disciplines, like bite-mark analysis, simply lacked a peer-reviewed scientific basis for core assumptions like the uniqueness of human dental patterns. Crime laboratories are undermined by unprofessional behavior and contamination. These and similar scandals come at great human cost. Innocent people are wrongly convicted and criminals remain free to harm more people. In addition, the weight and authority of science is party to act of injustice as the status of science is used to falsely convict innocent people in cases based on faulty science. This is a situation in need of serious reform.

The forensic science community is aware that there are problems in their practices and for the need of reform (Committee on Identifying the Needs of the Forensic Sciences Community, National Academy of Sciences 2009). Despite this awareness, the pace of reform progress is slow. One thing missing in this reform movement are philosophical perspectives on the nature of these scientific fields and their epistemic practices. In the paper, I draw on the work of contemporary philosophers of archaeology and use Allison Wylie’s concept of scaffolding to think about how to improve future forensic science (Chapman and Wylie 2016). The philosophical issues in archaeology and specifically the question of what is evidence and how do we make good inferences from it are directly applicable to the project of reforming the practices of forensic scientists. I argue that this body of knowledge offers a valuable perspective on the needed forensic reform process.

Chapman, R. and A. Wylie (2016) *Evidential Reasoning in Archaeology*. Bloomsbury Academic: New York.

Committee on Identifying the Needs of the Forensic Sciences Community, National Research Council (2009) *Strengthening Forensic Science in the United States: A Path Forward*. National Academies Press: Washington D.C.

Alignments – Mario Bunge’s General Black Box Theory and Contemporary Technoscience

Alfred Nordmann – Technische Universität Darmstadt, Germany

Hans-Jörg Rheinberger famously distinguished within experimental systems between technical objects and epistemic things. Whereas technical objects reliably play an assigned role in the design and execution of experiments, epistemic things stand at the center of scientific attention in that they pose a challenge to explanation and understanding. A black box serves as a technical object – and it is at issue whether it therefore ceased to be an epistemic thing. Of particular interest are therefore those black boxes which bracket or contain highly complex processes – black boxes, in other words, that cannot be opened to yield a more complete understanding but that need to stay shut in order to bypass otherwise insurmountable obstacles to knowledge.

For all kinds of black boxes, two questions can be asked – how do we know them and how do they embody or instantiate knowledge? The answers to these questions differ significantly, however, for those black boxes that rely on well-understood mechanisms and those that are systematically opaque and intellectually intractable. And yet, both types of black boxes are subject to a General Black Box Theory as articulated in 1962 by Mario Bunge.

The difference between the two types of black boxes and how we know them can be characterized with respect to a brief remark by Max Weber. The boxes that can be opened to reveal, upon closer inspection, a perfectly intelligible mechanism belong to a disenchanted world in which there is nothing magical and everything can be mastered through calculation. The boxes that behave predictably but, even upon closer inspection, appear to “work like magic” are technical objects that might as well be features of as yet largely uncomprehended nature. We learn to know them not by analyzing them intellectually but by learning to participate in their performances or behaviors – by becoming aligned or attuned to them.

It is here where it is worth to reconsider Mario Bunge’s General Black Box Theory. For Bunge, the Black Box is a general device by which to distinguish phenomenological and representational theories of knowledge: We know Black Boxes by studying their behavior which consists of the transformation of

an input into an output. For the purpose of a General Theory, black boxes are to be distinguished not by their representational content – e.g., the models of causal processes that are physically embodied by the Black Box – but by the different ways of transforming inputs to outputs. By the same token, the Black Box embodies and instantiates knowledge not necessarily by way of representations of causal relations, but by way of aligning input and output behaviors or by being attuned to the working order of natural and technical things – where this attunement can be achieved through technical optimization routines. The well-understood black boxes that rely on familiar mechanisms thus prove to be only a special case of the ones that are opaque and known only behaviorally or phenomenologically.

The Prodigal Genetics Returns

Aaron Novick – University of Pittsburgh, United States

The rise of evolutionary-developmental biology (evo-devo) provides an ideal opportunity to study the integration of scientific disciplines. Evolutionary theory, which descends from the modern synthesis of the 1930s-1950s, was formed largely without input from developmental biology. Evo-devo attempts to address this lacuna, integrating knowledge of development (especially developmental genetic) into evolutionary theorizing.

However, conceptual tensions arise during this process of integration. Ron Amundson, especially, has argued that key conceptual features of the theory of evolution that emerged during the modern synthesis preclude the relevance of developmental biology to that theory. Because of this, he claims, the modern synthetic theory and evo-devo are incommensurable. He traces the roots of this incommensurability to the separation of the study of heredity from the study of development by Mendelian geneticists. The synthetic evolutionary theory required input only from the study of heredity (i.e. Mendelian genetics), which fruitfully black-boxed developmental processes. Opening the black box, while informative about gene action during development, is irrelevant to evolutionary considerations, so long as evolutionary theory holds to the conceptual framework of the modern synthesis.

In this talk, I will challenge this understanding of the tensions between evo-devo and mainstream evolutionary theory. I will do so by reconsidering the nature of the split between heredity and development. My argument has three parts. First, I show that Mendelian geneticists (and, following them, the architects of the modern synthesis) understood this split quite differently than did developmental geneticists. Second, I show that Amundson's claims of incommensurability are justified only if one accepts the developmental

biologist's understanding of the split. Third, I show that contemporary evo-devo vindicates the Mendelians' understanding of the split. The conceptual tensions between evo-devo and mainstream evolutionary theory are therefore less than they appear.

More specifically, I argue that, for geneticists, the split between heredity and development was a methodological split. The methods for studying transmission genetics (sometimes identified as "heredity") were distinct from the methods for studying developmental genetics. Accordingly, each provided a different kind of understanding. Transmission genetics illuminated the correlations between genetic and phenotypic differences, while developmental genetics studied the processes by which genes produced their effects. Importantly, the Mendelians saw genes as playing crucial roles throughout the developmental process. It is in this sense that the split was methodological, not ontological.

Developmental biologists, by contrast, saw the split as ontological. The embryologist Frank Lillie, for instance, argued that there were principled reasons to think that genes could never explain core developmental processes. In this way, the study of heredity and the study of development were separated not just by their methods but by their ontologies. Insofar as the modern synthesis saw evolution as a fundamentally genetic process, then, it was conceptually separated from developmental biology. This is the source of Amundson's incommensurability. It exists, however, only from the side of the developmental biologists. Evo-devo, I argue, provides a developmental genetic theory that supports the Mendelian interpretation of the split, and so is not incommensurable with the synthetic theory.

Artifactualism and Philosophy of Science-in-Practice

Gui Sanches de Oliveira – University of Cincinnati, United States

The goal of this paper is to examine the promise of the artifactual view of models as a philosophical approach to science-in-practice. I point out shortcomings with existing formulations and motivate a more radical approach.

In traditional philosophy of science, models are analyzed as representations of real-world target phenomena, and the fundamental philosophical task is to solve the "problem of scientific representation" (Frigg 2003, 2006, Callender and Cohen 2006) and determine by virtue of what features models represent their targets. Influential accounts within the representational view have described the representational model-target relation as a matter of similarity (Giere 1988, 2010, Weisberg 2012) or isomorphism (van Fraassen 1980, 2008),

or with some deflationary, non-reductive definition (Suarez 2015, Morrison 2015).

Alongside these debates about representation, an exciting new philosophical view emerged that treats scientific models as tools, artifacts and instruments (see, e.g., Morrison and Morgan 1999, Knuuttila 2011, 2017, Isaac 2013, Currie forthcoming). The artifactual view suggests that models are not simply like ordinary tools in being useful for some end: models literally are concrete artifacts created by humans to enable specific forms of manipulation. For this reason, according to artifactualism, we cannot fully appreciate the role models play in advancing scientific knowledge until we see models as being on a par with other concrete instruments used in science.

Current formulations of the artifactual view incorporate many elements from the representationalist framework. Morrison and Morgan's (1999) seminal account explicitly assigns a representational function to models-as-instruments: "the model's representative power allows it (...) to teach us something about the thing it represents" (11). Knuuttila (2011) criticizes the representational approach for being too narrow in scope, but as an alternative she articulates an "artefactual approach to model-based representation" to elucidate "the actual representational means with which scientists go on representing" (263, italics added). Similarly, Currie (forthcoming) claims that engineering models do not fit the traditional representational view, but still he relies on thoroughly representational notions in his account of models-as-tools, such as the content-vehicle distinction.

These formulations of the artifactual view are compatible with representationalism because, I suggest, they are primarily concerned with the ontology of models. Ontologically speaking, the claim that models are 'artifacts' is perfectly compatible with the claim that models are 'representations': the categories are not mutually exclusive. But this ontological focus sells the artifactual view short. As I propose, reliance on the conceptual framework of representationalism brings the artifactual view back in line with the preoccupations of traditional philosophy of science and hampers adequate understanding of "science-in-practice" (Ankeny, Chang, Boumans, Boon 2011; Boumans and Leonelli 2013).

An alternative approach is to radicalize artifactualism, treating it not as an ontology of models but as a comprehensive framework for studying the construction and use of models-as-artifacts. Conceptually, radical artifactualism requires operationalizing traditional philosophical notions such as "similarity," "abstraction," and "idealization" non-representationally. And methodologically, it involves focusing on material rather than discursive aspects of scientific practice, an attitude inspired by the approach of cognitive archeology (Abramiuk 2012, Malafouris 2013).

Dag Prawitz. "The epistemic significance of valid inference". *Synthese*, 187(3):887– 898, 2012.

Machine Learning and Scientific Understanding

James Overton – Knocean Inc., Canada

Scientists increasingly rely on databases of published results. These databases take time and money to build and maintain, with expert curators being the slowest and most expensive part. As the volume of scientific data increases, expert curators do not "scale". Machine learning technologies promise to do the work faster and cheaper. Even if these technologies are effective, I argue that they may undermine scientific understanding.

The Immune Epitope Database (IEDB) contains more than a million experimental results from 19,000 published papers, covering nearly all the publications in immunology, allergy, and auto-immune research. The value of the database comes from its comprehensiveness and consistency, facilitating search and comparison. Moving data from heterogeneous publications to homogeneous database records requires curation: the paper must be read, understood, and translated into a standard format. Curation is far more than finding the measured values in the text. In order for the measurements to make sense, the details of the subjects and the experiment must be understood and encoded in a consistent way. Up to 400 fields are used to describe each experiment in the IEDB.

Initially, the IEDB attempted to use a large team of curators with undergraduate degrees. Each paper would be read by two curators, and a senior curator would resolve disagreements. This failed, and the IEDB switched to using a smaller team of PhD-level curators. An extensive curation manual has been developed, expressing the consensus of the IEDB curators on how a publications should be understood and entered into the database. The IEDB curation manual is shared on the Web as part of the IEDB documentation.

The average number of experiments per publication has increased much more quickly than the IEDB's budget for curators. It is tempting to face this challenge by automating the curation process, and the most promising candidates are machine learning systems.

The most common machine learning techniques today are statistical at root. Some of the core techniques are decades old, but newly viable because of cheap computations, vast information storage, and large amounts of curated data for training. Consider "System X", a hypothetical statistical machine learning system for IEDB curation. System X computes a function Y from publications to IEDB records. Function Y is not programmed into the system as

a set of rules -- it is learned through an iterative reinforcement process that compares thousands of publications to millions of curated IEDB records. The result is a "black box", a gigantic matrix of statistical associations, with no resemblance to the IEDB curation manual.

Philosophers of science have expressed concerns about just what statistical associations such as these can explain (e.g. Woodward 2003). I argue that System X is not intelligible on de Regt's (2017) Criterion for the Intelligibility of Theories. However modern machine learning techniques are diverse. I close with an example of how machine learning is actually used to assist IEDB curators, and a discussion of machine learning methods that may serve to enhance scientific understanding rather than undermine it.

Conceptual Schemes, Perspectivism, and Science for Neuroatypical Subjects

Themistoklis Pantazakos – University College London, United Kingdom

In this paper, I examine perspective not as a shift in a subject's point of view of the same object or phenomenon, but as a shift of conceptual scheme across subjects. I defend the notion of alternative conceptual schemes, and I draw tentative conclusions about scientific realism. I also examine what science for subjects of a different conceptual scheme means, and what it should mean. I do so by focusing on scientific practice regarding subjects with an Autism Spectrum Disorder (ASD).

I launch my investigation from Donald Davidson's famous contention that alternative conceptual schemes, and the very idea of a conceptual scheme, is nonsensical. Drawing on the relevant literature, I utilise three avenues to resist this conclusion.

First, I 'take away' a pillar sub-argument of Davidson's on which his main argument rests: that the existence of neutral content is necessary for the notion of a conceptual scheme (and the notion of alternative ones) to make sense, and that the existence of neutral content is philosophically unattainable. Contra Davidson, I submit that the notion of neutral content in the way Davidson means it – effable and theory-neutral – is indeed nonsense, but that alternative conceptual schemes do not need such a strict notion of neutral content to support their existence. I argue that the notion of the existence (not the detailed description) of an extra-linguistic, indescribable content is enough to do the ontological heavy lifting, and that this notion can be easily provided. Second, I showcase empirical work, which evinces the existence of conceptual schemes that have the same ambit and are importantly different between them (they consist alternatives). I bring to the fore Lajos Brons' concept of

‘applied relativism’, which defends the current existence of alternative conceptual schemes, and I develop some arguments of my own in its favour. Third, I anticipate an objection from the Davidsonian camp by arguing that, while applied relativism may not satisfy the Davidsonians regarding how deep differences run within extant alternative conceptual schemes. To do this, drawing mainly on Richard Rorty, I provide theoretical arguments for why the notion of a conceptual scheme is valid, and why we are philosophically legitimised to believe in the possible existence of alternative conceptual schemes, even radically different between them (with zero overlap of concepts).

After this work is done, I draw some tentative conclusions on what the above means for scientific realism, and I present the first results of my study regarding science made from subjects of a certain neurotypical conceptual scheme (psychiatrists and neurologists) for subjects of a highly atypical conceptual scheme (extreme cases of autism).

Scientific disagreement and explanatory relevance: a case study on the cholesterol - heart disease controversy

Veli-Pekka Parkkinen – University of Bergen, Norway

This presentation considers scientific disagreement about explanatory power and relevance. This may involve, but does not entail, the disagreeing parties endorsing mutually inconsistent theories. Criteria for evaluating explanatory relevance and power are typically learned tacitly, as a part of socialization to a research community. This makes such disagreements difficult to detect and easily confused with other types of disagreement. An illustrative case is presented from the history of research on hypercholesterolemia and heart disease.

The hypothesis that cholesterol causes atherosclerotic heart disease was first proposed in early 1900s, and soon became a topic of cross-disciplinary dispute. Up until the discovery of the low-density lipoprotein receptor in the 1970s, the evidence mostly consisted of pathological findings in humans and experimental animals, and epidemiological evidence that suggested that cholesterol is a difference-maker for population-wide heart disease prevalence. A notable party line between authority figures ran along the border between epidemiology and cardiology: some of the first to endorse the hypothesis were epidemiologists, while many of the most authoritative “cholesterol skeptics” were cardiologists.

Steinberg (2007) has suggested that cardiology’s focus on proximate explanantia of heart disease may partly explain their skepticism. Cardiologists

aim to uncover and intervene on the processes productive of the clinical manifestations of disease in an individual. For this purpose, cardiologists expect a satisfactory explanation to elaborate manipulable mechanisms. By contrast, epidemiologists track population-wide correlations between exposures and disease outcomes.

We start with Steinberg's suggestion, and argue that the skepticism within cardiology can be seen as a symptom of how mechanistic explanations work. The models of mechanisms proposed to explain phenomena do not describe the total, concrete process responsible for the explanandum phenomenon, but rather provide idealized templates that can be applied to different concrete cases regardless of variation in context-specific detail. Given a complex phenomenon such as heart disease, many different mechanistic models may be applicable. One of the most prominent contemporary mechanistic theories of atherosclerosis is the 'response to injury' hypothesis (Ross and Glomset, 1976a,b), which is an application of a mechanism of cell signaling to the growth of the atherosclerotic lesion. This mechanism does not consider cholesterol or other lipid metabolism. Thus, for the cardiologist who conceptualises atherosclerosis primarily as a runaway malfunction of cell-signaling, the epidemiological evidence for causality on the population level —even if true — would not fit the most relevant explanatory narrative.

We conclude by suggesting some neutral criteria of explanatory relevance and power in terms of a contrastive-counterfactual theory of explanation.

Ross, R. and Glomset, J. (1976a). The pathogenesis of atherosclerosis (first of two parts). *The New England journal of medicine*, 295(7):369–377.

Ross, R. and Glomset, J. (1976b). The pathogenesis of atherosclerosis (second of two parts). *The New England journal of medicine*, 295(8):420.

Steinberg, D. (2007). *The cholesterol wars: the skeptics vs. the preponderance of evidence*. Academic Press, San Diego, California, 1st ed edition.

Relations between psychotherapeutic practice and models of mental disorders

Julia Pfeiff – Leibniz Universität Hannover, Germany

In my talk, I will investigate how specific psychotherapeutic practices might have influenced the form of well-received explanatory models of mental disorders that have been put forward in clinical psychology.

More precisely, I will be concerned with an explanatory model of Major Depressive Disorder (MDD) that was proposed by the inventor of cognitive-behavioral therapy, Aaron Beck, in 2016. Using this model as an exemplar, I will

investigate whether (and if so, how) practical aims that psychotherapists have when explaining mental disorders to their patients within therapy might have influenced – and continue to influence – the content and form of such models. In a first step, I will point towards crucial differences between Beck's (2016) explanatory model of MDD and typical explanatory models of mental disorders from fields, such as psychiatry. Among the features which are specific for his model are the presence of folk-psychological vocabulary as well as the normalization (compare Bolton, 2008) of the patient's experience.

Having thereby set the stage, I will argue for my main thesis, namely, that these very features are due to the influence of several practical aims that arise within explanatory practices in the context of psychotherapy.

To do so, it will be necessary to first identify those practical aims. I approach this question through the use of data from a qualitative interview study on explanatory practices within clinical psychology. On this basis, I will conclude that several such goals arise, for example, shifting the blame away from the patient, motivating her to undergo structured treatment and enabling her to cope better with her symptoms.

Secondly, I will examine the central explanatory strategy of Beck's model of MDD and suggest that it aligns well with the practical aims I identified before.

Thirdly, I will argue that this influence from explanatory practices within psychotherapy on the structure and content of the explanatory model can be understood as due to the fact that Beck's model is primarily based on observations from clinical practice, including patient's self-reports.

Finally, I will examine the extent to which this finding can be generalized to other explanatory models in clinical psychology and I will discuss potential implications of my thesis that specific aims arising within psychotherapeutic practice exert substantial influence on the content and structure of explanatory models of mental disorders.

Contextualising values in science: simplicity, completeness and carefulness

Karoliina Pulkkinen – University of Cambridge, United Kingdom

In this paper, I apply historical methods to philosophical notions on values in science. I examine the priority dispute concerning the discovery of the periodic system, and argue that the differences between the classifications of J.A.R. Newlands, Julius Lothar Meyer, and Dmitri Mendeleev can be explained in terms of three diverging values: simplicity, completeness and carefulness. Where Newlands stressed having identified a “simple relation”, Lothar Meyer warned against assuming any underlying simple laws. Where Newlands and

Mendeleev included many elements, Meyer included fewer. While Meyer only wanted to include “well-characterised” groups of elements, Mendeleev discussed little-known elements in his arrangement.

In the philosophical literature, characteristics like simplicity, completeness and carefulness have usually been referred to as epistemic or cognitive values. Thomas Kuhn, by popularising the terminology of epistemic values, meant them to provide “the shared basis for theory choice” (1977, 322). While there is a flexibility in interpreting values, Kuhn maintained that they could be discussed across different historical contexts. This universality was questioned especially by Phyllis Rooney (1992) and Helen Longino (1996). Rooney argued against universal epistemic values on the grounds that features that we would deny the epistemic status were considered epistemic in the past. In light of this, Longino argued that we should adopt a more local conception of values (Longino and Lennon 1997).

This paper develops Longino’s suggestion of local values by demonstrating that definitions of carefulness, simplicity, and completeness can be formulated to reflect the local context where they arise from. First, I show how Hasok Chang’s definition of completeness (2012, 22) can be altered to fit the context of classifying the chemical elements. I then argue that we need analyses of simplicity independent of ontological simplicity (parsimony). As Newlands stressed the “simple relation” expressed by his Law of Octaves, but the number of elements was growing, alternative analyses of syntactic simplicity are needed. Finally, I introduce a new category of carefulness to account for Meyer’s approach to classification, as related categories of accuracy, precision, and exactness are not as effective in capturing Meyer’s approach.

Bringing together historical and philosophical methods in this particular case suggests that it is sometimes inappropriate to categorise some value as purely pragmatic, epistemic, aesthetic, or cognitive. Instead, this historical case-study suggests taking some values as potential compounds of different functions. In order to illustrate this claim, I argue that for Mendeleev completeness was simultaneously a pragmatic and an epistemic value. This kind of approach of localising values, and bringing forth to their complexity, allows us to see how values reflect historical contexts. Global accounts of values in science would benefit from this approach, as contextualising values would allow us to inspect, and to compare and contrast, different hierarchies of values in different historical contexts.

Chang, Hasok. 2012. *Is Water H₂O?* Dordrecht: Springer.

Kuhn, Thomas. 1977. ‘Objectivity, Value Judgment, and Theory Choice’. In *Essential Tension*, 320–39. Chicago: University of Chicago Press.

Longino, Helen E. 1996. ‘Cognitive and Non-Cognitive Values in Science: Rethinking the Dichotomy’. In *Feminism, Science, and the Philosophy of*

Science, edited by Lynn Hankinson Nelsin and Jack Nelson, 39–58. Dordrecht: Kluwer Academic Publishers.

Longino, Helen E., and Kathleen Lennon. 1997. 'Feminist Epistemology as a Local Epistemology'. *Proceedings of the Aristotelian Society, Supplementary Volumes* 71: 19-35-37–54.

Rooney, Phyllis. 1992. 'On Values in Science: Is the Epistemic/Non-Epistemic Distinction Useful'. *PSA 1992: Proceedings of the 1992 Biennial Meeting of the Philosophy of Science Association* 2: 13–22.

How rational choice-theoretic models in economics can causally explain

Beau Revlett – University of Kentucky, United States

Recently, Julian Reiss argued that we do not understand how economic models can explain, and Anna Alexandrova and Robert Northcott argued that rational choice-theoretic models in economics cannot explain. In this paper, I argue for a reinterpretation of Alexandrova and Northcott's example of the design of the FCC Spectrum Auctions. I argue that, given my reinterpretation, rational choice-theoretic models in economics can ground capacity claims, can count as explanatory on James Woodward's account of causal explanation, and can reveal the role rationality plays in complex decision-making situations that humans face. I start with the latter, which grounds the former two.

One way of seeing rational choice theory is as an attempt to answer the question, "If humans were purely rational, what would they do?" Of course, humans are not purely rational, given rational choice theory's definition of purely rational as always maximizing expected utility, where expected utility is defined as usual. Thus, rational choice-theoretic models cannot explain all human decisions.

Nonetheless, as Alexandrova and Northcott have described well, the FCC successfully designed an institution with significant aid from rational choice-theoretic models from economics. The success of the institution depended on the designers' ability to predict how humans would behave when interacting with the institution. Thus, this is an example of rational choice-theoretic models in economics aiding in grounding successful predictions of human behavior, despite the agents of rational choice theory being purely rational.

How did the FCC do it? In general, to successfully design an institution based on rational choice-theoretic models, one must predict in what ways humans will fail to maximize expected utility during interaction with the institution. Since the agents of rational choice theory are purely rational, such prediction is beyond the scope of rational choice theory. Alexandrova and Northcott

recognize this. They point out that to predict how humans would actually behave, the FCC hired experimental economists.

From this, Alexandrova and Northcott concluded that rational choice-theoretic models in economics do not explain; instead, the experiments they inspire do. I reject Alexandrova and Northcott's conclusion. I claim that rational choice-theoretic models in economics do explain why people make the decisions that the models predict they will. What the experiments show is that humans do not always make such decisions, and what effect this will have on interaction the FCC Spectrum Auctions. I will discuss examples of models and practices that contributed to the design of the FCC Spectrum Auctions. To clarify what the examples show about explanation, I will connect my account to Nancy Cartwright's theory of capacities and Alisa Bokulich and James Woodward's theories of explanation.

The scope of evolutionary explanations as a matter of “ontology-fitting” in investigative practices

Thomas Reydon – Leibniz Universität Hannover, Germany

Both in academic and in public contexts the notion of evolution is often used in an overly loose sense. Besides biological evolution, there is talk of the evolution of societies, cities, languages, firms, industries, economies, technical artifacts, car models, clothing fashions, science, technology, the universe, and so on. While in many of these cases (especially in the public domain) the notion of evolution is merely used in a metaphorical way, in some cases it is meant more literally as the claim that evolutionary processes similar to biological evolution occur in a particular area of investigation, such that full-fledged evolutionary explanations can be given for the phenomena under study.

Such practices of “theory transfer” (as sociologist Renate Mayntz called it) from one scientific domain to others, however, raises the question exactly how much can be explained by applying an evolutionary framework to cases outside the biological realm. Can applications of evolutionary theory outside biology, for example to explain the diversity and properties of firms in a particular branch of industry, of institutions in societies, or of technical artifacts, have a similar explanatory force as evolutionary theory has in biology? Proponents of so-called “Generalized Darwinism” (e.g., Aldrich et al., 2008; Hodgson & Knudsen, 2010) think it can. Moreover, they think evolutionary thinking can perform a unifying role in the sciences by bringing a wide variety of phenomena under one explanatory framework.

I will critically examine this view by treating it as a question about the ontology of evolutionary phenomena. My claim is that practices of applying evolutionary thinking in non-biological areas of work can be understood as what I call “ontology-fitting” practices. For an explanation of a particular phenomenon to be a genuinely evolutionary explanation, the explanandum’s ontology must match the basic ontology of evolutionary phenomena in the biological realm. This raises the question what elements this latter ontology consists of. But there is no single answer to this question – there is ongoing discussion about what the basic elements in the ontology of biological evolutionary phenomena and explanations are and how these are to be conceived of. (Consider for example the debates on the nature of selection and of the units of selection.) Therefore, practitioners from non-biological areas of work cannot simply take a ready-for-use ontological framework from the biological sciences and fit their phenomena into it. Rather, they tend to pick those elements from the biological evolutionary framework that seem to fit their phenomena, disregard other elements, and try to construct a framework that is specific to the phenomena under study. By examining cases of such “ontology fitting” we can achieve more clarity about the requirements for using evolutionary thinking to explain non-biological phenomena. I will illustrate this by looking at an unsuccessful case of “ontology fitting” in organizational sociology.

Aldrich, H.A., Hodgson, G.M., Hull, D.L., Knudsen, T., Mokyr, J. & Vanberg, V.J. (2008): ‘In defence of generalized darwinism’, *Journal of Evolutionary Economics* 18: 577-596.

Hodgson, G.M. & Knudsen, T. (2010): *Darwin’s Conjecture: The Search for General Principles of Social and Economic Evolution*, Chicago & London: University of Chicago Press.

Mathematisation: a pragmatist account

Davide Rizza – University of East Anglia, United Kingdom

Recent work in philosophy of mathematics, especially Lange, Pincock, Baker made significant progress in identifying subtle features of the function of mathematical resources in scientific enquiry, with a definite emphasis on explanation.

These analyses, however, are framed in connection with metaphysical problems: Lange talks about distinctively mathematical explanation in terms of modal strength of law-like connection; Baker talks about the benefits deriving from topic generality in mathematical explanation and ties them to mathematical realism; Pincock studies abstract explanation and is involved in a

dependency problem of abstract features and empirical ones. In all cases, an important dimension of the explanatory function performed by mathematisation is singled out, but it is attached to metaphysical categories that ground its significance or affect its nature and yet remain alien to the needs and motivations of scientific practice. This drawback can be avoided by offering an account of mathematisation and of mathematical explanation along pragmatist lines. One such account preserves and extends all of the insights generated by the recent literature without allying them to traditional metaphysical stances. Its background is the conception of enquiry developed by John Dewey. I propose, after Dewey's contributions, to regard mathematisation as a particular type of enquiry or a particular way of conducting empirical enquiry, which produces distinctive secondary objects (i.e. selective analyses of given experience that allow intelligent intervention upon it). The secondary objects of mathematisation are methods of focussing attention on the features of a problem that can then be subjected to particular forms of reasoning or that open the problem to the application of further mathematical techniques. A mathematised problem is in the first instance an environment that allows formal experimentation regimented by proof and, in the second instance, a basin of attraction for external mathematical instruments and techniques.

Against the background of this conception of mathematisation, its success amounts to the wealth of connections it establishes between disparate empirical contents and to the solidity of the means of control it generates on the problems of interest. Then 'mathematical explanation' is to be understood as a generic expression to indicate that a significant measure of success is achieved. A problem is mathematically explained when a mathematical model is offered for it that allows a more transparent and tighter management of the consequences produced by the setup from which that problem arose. An explanation is distinctively mathematical because this can sometimes be done without extensive accompanying experimental research. The generality of a mathematical explanation amounts to the fact that its formal constraints can be deployed as effective instruments in a variety of empirical contexts. Its abstract character consists in detachment from direct reference to conditions present within a fixed empirical environment. Abstract character could only lead to dependency problems if the mathematical instrument was reified into an object of knowledge, as opposed to being recognised in its role as a function of enquiry. In short, the insights of Baker, Lange and Pincock about topic generality, distinctive mathematical character, and abstract character respectively can be captured in a way that remains close to the practice of modelling and its demands, without any loss of depth in the appreciation of scientific enquiry or unduly intrusions of traditional metaphysical concerns.

The Citizen Science Movement According to Feyerabend

Sarah Roe – Southern Connecticut State University, United States

The slogan 'anything goes' first appears in Paul Feyerabend's book *Against Method*. Many have speculated on what exactly was meant by the slogan and even more philosophers and scientists have quickly discarded Feyerabend's antidote as the obvious ramblings of a madman. Instead of discarding Feyerabend, I utilize his work to better understand the new citizen scientist movement.

Citizen scientists partake in scientific discovery, monitoring, data collection, and experimentation across a wide range of scientific disciplines. The information collected by citizens is most often used to better understand and predict phenomena like climate change, overexploitation, invasive species, land use change, pollution, etc. Citizens are able to collect fine-grain data over regional, and sometimes continental, extents and decadal time scales, something professional scientists are unable to accomplish alone. For instance, citizen science data bases, like the one geared toward bird watchers, eBird, have more than 100 million records on file. As such, it is estimated that the scientific benefit of citizen scientists easily amounts to millions in in-kind economic worth. It is clear citizen scientists may offer increasing support for science.

For Feyerabend, a healthy society needs active citizen participation in science. The modern citizen science movement would not have surprised Feyerabend, rather it would have excited him! To be sure, the citizen science movement has grown naturally as a societal feedback mechanism between citizens and experts of science. If my argument holds, Feyerabend would have championed this as progress, movement toward a more equal balance between the importance of citizens and societal aims and the value of the scientific process. However, I argue that Feyerabend would champion a more radicalized citizen science, one that allows for the possibility of integrating citizens into every level of the scientific process. I delineate the five goals set out by the citizen science movement and argue that, currently, only two of the goals are being promoted by the movement. Feyerabend taught us that while the current citizen science movement is primarily focused on what the citizen can do for science and what the citizen can learn from science, the movement, moving forward, should also focus on what science can do for the citizen and what science can learn from the citizen. I hope my arguments provide a better understanding of how the citizen science movement can best promote scientific education and a broader knowledge to participants, increase citizen interest in conservation and policy, increase both citizen local and national

engagement, while promoting an awarding experience for both the expert and laymen alike.

By understanding the contemporary citizen science movement alongside Feyerabend's views, the way forward becomes clearer. In order for science and society to benefit in the ways the citizen science movement hopes, more integration between citizens and science must occur, and importantly, at every level.

The Division of Replication Labor

Felipe Romero – University of Groningen, Netherlands

Science is going through a worrisome replication crisis—a crisis that arose from an explosion of replication failures of published scientific studies (ranging from psychology studies to medically critical studies in cancer research), and constitutes the most serious epistemological and ethical problem for science today (Open Science Collaboration, 2015). To increase replicability we encounter three major problems. The first problem is epistemic: we need independent replications. Without them, we cannot rule out systematic error and increase our confidence in published findings. The second problem is sociologic: replication failures often lead to stalemates. Many failed replication reports are perceived as hostile attacks, and original researchers and replicators question each other's qualifications and intentions instead of engaging in fruitful discussions about ideas. The third problem is economic: there are few material incentives to replicate. Despite the recent acknowledgment of the epistemic importance of replicating, replication is still under-rewarded second-class work.

There are various proposals that suggest intervening on the social structure of science to address these problems. These proposals specify the appropriate roles, responsibilities, and communication protocols for scientists who conduct replication research. Inspired by the philosophical literature on the division of cognitive labor (Kitcher, 1990), I call these proposals replication labor schemes. The most discussed schemes proposals are:

- (1) the proposer scheme (Roediger, 2012; Cesario, 2014), which requires proposers of findings to ensure the replicability of their own findings;
- (2) the consumer scheme (Pashler and Harris 2012), which requires researchers to ensure the replicability of crucial findings they rely on;
- (3) the students scheme (Frank and Saxe, 2012; Standing et al., 2014), which requires students in methods classes to attempt replications as part of their training; and

(4) the multi-site scheme (Open Science Collaboration, 2012), which requires various laboratories to coordinate replication attempts in their fields.

In my talk, I identify the limitations of these schemes and propose a different solution. I join those philosophers who think that modern science requires dividing cognitive labor because the complexity of modern science exceeds the cognitive capacities of individuals (Weisberg and Muldoon, 2009). Contrary to what we see in practice, I argue that replication labor also requires such a division.

My argument proceeds in three steps. First, I present evaluation criteria that replication labor schemes should satisfy to solve the epistemic, sociologic, and economic problems. These criteria require making replication labor outcome-independent (i.e., replicators should not have conflicting interests regarding replication outcomes), systematic (i.e., crucial findings should be replicated multiple times), and sustainable (i.e., replication labor should be a standard practice). Second, using these criteria, I evaluate the four extant scheme proposals and defend my proposal, which I call the professional scheme. This scheme recommends creating a distinct reward system for scientists whose primary function is to do confirmation/replication labor (i.e., replication, reproduction, meta-analysis, and theory criticism). To conclude, I argue that one way of implementing the professional scheme is establishing confirmation research tracks and I discuss how different stakeholders involved in the research process (from funding agencies to journals) should support such a system.

Epistemic Risk, Scientific Significance, and Conceptual Normativity

Joseph Rouse – Wesleyan University, United States

Recent philosophical work revives and extends earlier discussions of “inductive risk” in hypothesis acceptance as appropriately raising non-epistemic normative considerations within scientific practice. Douglas (2000) extends inductive risk beyond hypothesis acceptance to include choices of methodology, evidence descriptions, and data interpretation. Biddle (2016) expands Douglas’s reasoning further, incorporating narrowly inductive risk within the broader category of “epistemic risk,” because errors can arise in accepting or rejecting methodologies, background assumptions, policies, and more.

I highlight two related considerations that further extend and re-characterize these expansions of the normative concerns appropriately arising within scientific practice. The first recognizes that we assess scientific claims and projects for scientific significance as well as correctness. Some research

projects would not matter scientifically even if their outcomes were determined correctly. Assessments of significance can go awry not only in whether projects and claims are important, but also in how and why so. Two examples are illustrative. The 1914 chemistry Nobel-winning research that precisely determined many elements' atomic weights presumed that these values provide insight into atomic structure rather than the relative preponderance of different isotopes. The successful sequencing of human and other genomes, while a very important achievement, also showed that knowledge of genome sequences has different significance than was originally put forward as rationales for the project's requisite reorientation of scientific priorities. The normative considerations bearing upon the articulation and assessment of scientific significance extend well beyond epistemic norms.

The second consideration involves the broader inferential roles of scientific claims, models, and concepts. The inductive/epistemic risk literature emphasizes the direct practical consequences of errors in accepting hypotheses, methodologies, or definitions: misdiagnosed disease, deleterious environmental consequences, ineffective or destructive military strategies, and other unsuccessful or harmful applications of scientific claims and methods. The integration of scientific concepts and methods with broader cultural or political concerns extends further than technological applications and public policy choices, however, and requires assessment within a broader normative register. A rich anthropological, historical, and feminist-theoretical literature explores how these broader conceptual concerns are important for adequate critical engagement with scientific practice (classic examples include Haraway 1989 on field primatology, Reardon 2005 on evolutionary genetics, or Martin 1994 on immunology). These considerations cannot be reduced to epistemic risks, since the inferential patterns involved are non-monotonic, and hence cannot readily be captured as explicit conditional claims. They also cannot be relegated to considerations external to the sciences proper, because these inferential relations are bi-directional, and often contribute extensively to scientific significance, and guide scientific practice. More broadly, these considerations contribute to the recognition that the normative significance of scientific conceptualizations extends well beyond epistemic assessments of errors in scientific knowledge, methodology, or their applications.

Biddle, Justin 2016. Inductive Risk, Epistemic Risk, and Overdiagnosis of Disease. *Perspectives on Science* 24:192-205.

Douglas, Heather 2000. Inductive Risk and Values in Science. *Philosophy of Science* 67:559-79.

Haraway, Donna 1989. *Primate Visions*. New York: Routledge.

Martin, Emily 1994. *Flexible Bodies*. Boston: Beacon.

Reardon, Jenny 2005. *Race to the Finish*. Princeton: Princeton University Press.

The Biochemical Roots of the pH Value and the Glass Electrode

Klaus Ruthenberg – Coburg University of Applied Sciences and Arts, Germany
Hasok Chang – University of Cambridge, United Kingdom

Acidity is a chemical concept which usually is described and discussed with emphasis on theoretical interpretation. In contrast, the present paper tells a parallel story of scientific motivations and practices. It explores the early history of the pH concept and of the glass electrode. Both are bioanalytical achievements, that is, both were developed out of biochemical motivations. In 1906, the physiologist Max Cremer published the results of his attempts to mimic biological membranes and their electrophysiological effects. He adopted a specific Daniell cell Hermann von Helmholtz and his co-worker Wilhelm Giese had tried out for the purpose to study electric potentials more than 20 years before without recognizing the selectivity of their device. This selectivity – the impact of aqueous acidity on glass surfaces – was the main subject of Cremers study, although he had no appropriate idea or theory about it. He used the well developed skills of the glassblowers of his time and applied small and thin glass bulbs to construct the first “glass cell” to measure acidity.

Three years later, the Danish biochemist Søren Sørensen, working for the Carlsberg brewery in Copenhagen and being interested in the kinetics of enzymatic cleavage, suggested the algebraic redraft of an already well-known empirical fact, namely the concentrations of the hydroxide- and the hydrogen-ions in water. His main idea was to describe the influence of acidity and alkalinity on these enzymatic reactions by a classification number. That redraft proposed to call the negative logarithm of the molar hydrogen-ion concentration “hydrogen exponent”. Hence, Sørensen invented the pH value, which very soon made its way to the whole of the chemical sciences and is now one of the globally best known (but not equally well understood) expressions from chemistry. Moreover, it is doubtless the most frequently determined value in this science and its neighbors. It is nevertheless still mysterious why this numerically disguised measured value has been so extraordinarily successful. To illustrate that point: This concept would arguably not have been invented any more after the development of electronic calculators. Perhaps just its assumed (but – at least for students – as well dangerous) simpleness and wide application in the biomedical and ecological fields is the key to that mystery.

Although it looks like a direct reading method of acidity, the application of glass-electrodes for pH-measurements relies on a kind of an arteficially brought about chemical reaction, and is much more restricted than expected in the first place.

Cremer, M. (1906): Über die Ursache der elektromotorischen Eigenschaften der Gewebe, zugleich ein Beitrag zur Lehre von den polyphasischen Elektrolytketten. *Zeitschrift für Biologie* 47, 562-608.

De Levie, R. (2014): A pH centenary. *Electrochimica Acta* 135, 604-639.

Dole, M. (1941) *The Glass Electrode*. New York: John Wiley & Sons.

Ruthenberg, K. & Chang, H. (2017): Acidity: Modes of characterization and quantification. *Studies in History and Philosophy of Science* 65-66, 121-131.

Scholz, F. (2011): From the Leiden jar to the discovery of the glass electrode by Max Cremer. *Journal of Solid State Electrochemistry* 15, 5-14.

Sörensen, S. (1909): Enzymstudien II. *Biochemische Zeitschrift* 21,131-200.

Scale Models in Civil Engineering: Creative Similarity in Representational Practices

Julia Sanchez-Dorado – University College London, United Kingdom

Scale models have been used in hydraulics and civil engineering at least since the 17th century, when Galileo started experimenting with small vessels in water tanks. Ever since they have served multiple purposes, from research in fluid dynamics, to the prediction of floods, and the construction of ships, bridges, and dikes.

However, in contemporary philosophy of science, scale models are often disregarded in favour of mathematical or computational models. This is manifest, first, in the little attention given to scale models in the literature on scientific representation and explanation, with exceptions like Pincock (2012), Sterrett (2006; 2017), and Weisberg (2013). Second, it is frequently assumed that scale models have a lower epistemic status than other types of models, and for this reason taken as mere illustrative or pedagogical tools and compared with toys and collector's miniatures (Frigg & Hartmann 2012; critique in Sterrett 2017). Third, scale models are usually identified with instruments of a past era, when resources such as computational models did not exist (Oreskes 2007; reply in Sterrett 2017).

The aim of this paper is to cast doubts on the aforementioned assumptions about the interest of scale models for philosophers of science. To do so, I argue that an approach in integrated HPS, which looks into historical episodes in civil engineering, can be particularly insightful to discuss the epistemic value of scale models.

I use the documented history of the foundation of the Waterways Experimental Station in the U.S. (WES) as a case in point. Before and after the foundation of the WES in 1929, there were fruitful disagreements within the

U.S. Army Corps of Engineers about the kind of knowledge that scale models could provide in comparison with numerical – and later on, computational – models (Robinson 1992). From here, I conclude that scale models have been historically used in civil engineering for their demonstrated predictive and explanatory power. In other words, they are neither mere illustrative tools nor (totally) replaceable for numerical or computational models.

Then, I look more specifically into historical reports, interviews, and technical manuals of the construction of the Mississippi Basin Model (1943-1970s). Here we can find evidence of the particular actions and judgments that engineers developed during the practice of designing, building, and evaluating this model (Foster 1971; Rouse 1950; U.S. Army Corps of Engineers 1997). I conclude that debates in philosophy of science can be substantially enriched by taking seriously into consideration the content of these reports and manuals, particularly in relation to the philosophical problem of scientific representation (as discussed by Weisberg 2013; van Fraassen 2008). In this, history and philosophy of science are brought one step closer.

Frigg, R. and Hartmann, S. (2017) “Models in Science”. The Stanford Encyclopedia of Philosophy (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2017/entries/models-science/>

Foster, J. E. (1971) Mississippi Basin Model Report 1-6. History and Description of the Mississippi Basin Model. US Army Engineer Waterways Experiment Station, Vicksburg, Mississippi.

Oreskes, N. (2007) “From Scaling to Simulation: Changing meanings and ambitions of models in geology”. In Angela N. H. Creager, Elizabeth Lunbeck, and M. Norton Wise (Eds.): *Science Without Laws*. Duke University Press, Durham, London.

Pincock, C. (2012) *Mathematics and Scientific Representation*. New York: Oxford University Press.

Robinson, M. C. (1992) “Rivers in miniature: The Mississippi Basin Model”. In B. W. FOWLE (Ed.), *Builders and fighters: U.S. Army Engineers in World War II*. Fort Belvoir, VA: Office of History, United States Army Corps of Engineers.

ROUSE, H. (1950): *Engineering Hydraulics*. Proceedings of the Fourth Hydraulics Conference. Iowa Institute of Hydraulic Research, 1949.

Sterrett, S. G. (2006) “Models of Machines and Models of Phenomena”. *International Studies in the Philosophy of Science*, 20:01, 69-80.

Sterrett, S. G. (2017) “Experimentation on Analogue Models”. In L. Magnani and T. Bertolotti (eds.) *Springer Handbook of Model-Based Science* (pp. 857-878). Dordrecht: Springer.

U.S. Army Corps of Engineers (1997) "Mississippi Basin Model at Jackson." Water Resources: Hydraulics and Hydrology: Interview with Margaret S. Petersen. USACE Engineering Pamphlets EP 870-1-60.

van Fraassen, B. (2008). *Scientific Representation*. Oxford University Press.

Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. New York: Oxford University Press.

Conceptions of experimental control in 19th-century life sciences

Jutta Schickore – Indiana University, United States

Arguably, controls are a key feature of scientific experimentation. However, there are very few systematic studies of concept of experimental control. Historical and philosophical analyses of control experiments have mostly concentrated on randomized controlled trials as "gold standard" for experiment-based inferences in medicine and specifically on the concept of randomization (e.g. Hacking 1988, Cartwright 2010). But apart from a paper by Edwin Boring on the nature and history of experimental controls (Boring 1954), broader systematic analyses of the concept of control, the epistemological significance of the practice of controlling, and the conditions of the emergence of the methodological ideas behind experimental controls do not exist. This paper offers contributions to such an analysis of experimental controls. The focus is on the life sciences in the German lands in the first half of the 19th century.

My contribution builds on Boring's 1954 analysis. Boring helpfully distinguishes among three meanings of "control" in the context of experimentation: "control" in the sense of restraint (keeping conditions constant); "control" in the sense of guided manipulation (causing an independent variable to vary in a specific manner); and "control" in the most general sense of check or comparison. He also suggests that the first philosophical conceptualization of controlled experiments can be found in John Stuart Mill's *System of Logic* and that the very term "control" appears in the scientific literature only in the late 19th century.

Focusing on working scientists' own conceptualizations of experimental practice and of methodological strategies of experimentation, I highlight the dramatic difference between the pragmatic concerns of the experimenters and the systematic concerns of philosophers. The practitioners often (not always!) perceived Mill's conception of experimental methods as unhelpful and advanced their own conceptualizations of controls, causes, and complexity in experimentation. I examine examples of practitioners' conceptions of experimental control, showing that the term "control" did appear in the early

19th-century scientific literature, namely in the context of experiments on plant growth and plant nutrition in the German lands. Control experiments [Controlversuche] were a part of agricultural field trials to assess the influence of air, water, minerals and organic materials on plant growth. I analyze the meaning of this conception of control and put my analysis in historical perspective, showing that there are two related contexts from which this notion of control emerged: early 19th-century socio-political concerns about controlling populations and methodological discussions in the emerging organic chemistry around 1800. My analysis of these early instances of “control” suggests a systematic distinction between two kinds of check and comparison that is not represented in Boring’s threefold distinction of “control” but is well represented in late 19th-century discussions of experimental methods in the life sciences.

Boring, Edwin G. 1954. The Nature and History of Experimental Control. *The American Journal of Psychology* 67, 573-589.

Cartwright, Nancy. 2010. What Are Randomised Controlled Trials Good For? *Philosophical Studies* 147, 59-70

Hacking, Ian. 1988. Telepathy: Origins of Randomization in Experimental Design. *Isis*, 79: 427-51.

Bridging the gap between populations and individuals in personalized medicine

Raphael Scholl – University of Cambridge, United Kingdom

The goal of personalized medicine is to stratify patient populations into subgroups according to biologically relevant individual variations. In principle, these variations could be in lifestyle or environment; in practice, they are usually genetic. The hope is that subgroups will exhibit meaningful regularities that are directly relevant to individual patients. We may be able to explain why a risk exists or a disease develops in members of a particular subgroup; what course of disease that subgroup should expect; or how the subgroup will respond to different kinds of therapies.

However, this project has turned out to be more challenging than early proponents expected. Around the time of the completion of the human genome project, it was expected that association studies would find a handful of genetic variations with relatively large effects that are relevant to explanation, prognosis, and therapy. But most of our data indicates that medically relevant genetic variations are for the most part rare and heterogeneous. This presents an unexpected epistemological challenge.

Correlational studies, such as genome-wide association studies, are often insufficient for investigating causal structures in which the same effect can be produced by a large range of different causes, where each cause occurs only infrequently, and where each cause typically only has a small effect size.

In this talk I will take a close look at some of the practices that allow scientists to overcome the limitations of traditional methodologies. In particular, I will focus on a case study from recent cancer research. Over the past decade, the genetic causes of tumors have been found to exemplify the epistemic challenge of personalized medicine: they are both extremely rare and strikingly heterogeneous. At the molecular level, the somatic mutations of most tumors differ from those of most other tumors, even when their clinical phenotypes are indistinguishable. Population-based correlational methods are usually insufficient to attach reliable regularities to such rare variants.

In recent years, however, progress has been made by a number of related techniques that can be grouped under the label of "network-based stratification". The principle of these techniques is the following: While it is often difficult or impossible to attach regularities to rare and heterogeneous gene variants themselves, the sought-after regularities can be discerned more readily at the level of **gene network modules**. A tumor's prognosis and response to therapy seems to depend significantly on which network modules are altered or disrupted by individual mutations. Thus, network reconstruction promises to deliver information about the causation, prognosis, and response to therapy of individual cancers in individual patients.

In conclusion, I will highlight some further practices that promise to bridge the gap between populations and individuals, such as the often maligned n-of-1 trial. I will argue that medical epistemology cannot meet the challenges of personalized medicine unless it goes beyond its current focus on population-based correlational studies.

Values in Science, Public Trust, and Transparency

Andrew Schroeder – Claremont McKenna College, United States

There is a growing (and I think clearly correct) consensus among philosophers of science that science can't be value-free or "objective" in the way most people think. Instead, core parts of the scientific process — classifying data, evaluating hypotheses, choosing an experimental design, and/or analyzing and reporting results — require an appeal to non-epistemic values.

We have empirical evidence, however, that when members of the public believe that scientific results depend on or were influenced by a scientist's values, that undermines the trust they have in those results. I argue that this

response is rationally justified: on the basis of case studies from epidemiology and climate science, I show how key conclusions in major scientific studies can crucially depend on underlying value judgments. Unfortunately, there is often no way for non-specialists to determine, concerning some particular result, whether or not this has been the case. And, importantly, in such cases there are few generally-accepted principles of scientific ethics which significantly constrain scientists' value judgments. Thus, rejecting the value-free ideal for science both **does** and in many important cases **should** undermine public trust in science.

What is the solution to this problem? Several authors have proposed transparency: scientists should clearly identify and describe the value judgments they rely on in the course of their work. I argue that this is insufficient. On the descriptive side, we have empirical evidence showing that reporting values does not increase public trust in science. (In fact, it decreases it.) And, given the highly complex and technical impact of many value judgments on scientific results, I argue that transparency also probably should not increase public trust in science, at least in many important cases. Non-specialists usually have no ability to “reverse engineer” a scientist's value judgments, to determine what the results would have been under some alternative value judgment. Transparency is therefore insufficient to restore public trust. (It is for similar reasons that many medical schools and journals have moved away from simply requiring transparency about conflicts of interest, and are instead attempting to prevent such conflicts from arising at all.)

I conclude by suggesting an alternative. The real problem, I argue, is not that scientists' value judgments are hidden. (Transparency would address this.) Instead, it is that scientists' value judgments are (relatively) unconstrained. I therefore propose adopting (a more nuanced version of) the following as a principle of scientific ethics: when scientists must make value judgments in the course of their research, they should (in many cases) not rely on their own values. They should instead ground their work in democratic values — roughly, the values of the public or its representatives. Adopting and enforcing this principle, I argue, rationally should increase public trust in science. Unfortunately, we don't yet have any evidence about whether it will increase public trust in science, but I think there are reasons for optimism.

Competition and the Creation of Neuroimaging: The History of Positron Emission Tomography 1976-1985

Rick Shang – Washington University in St. Louis, United States

Competition is one of the fundamental elements of scientific development. But how does competition shape scientific development? Few philosophers attempt to give an answer. Lakatos, for example, endorses a “winner-takes-all” account, according to which, the winner is better than the loser according to all or most of the relevant evaluative criteria and replaces the loser (Lakatos 1970).

The history of neuroimaging defies such a “winner-takes-all” account. In 1974, researchers at Washington University in St. Louis created a prototype of a new medical imaging device called the Positron Emission Transaxial Tomograph (PETT) scanner. By 1977, fewer than 20 articles a year were published about research that used the most advanced PETT scanner. Many researchers wondered whether the marginal increase in precision or accuracy over competitors could justify the million-dollar price tag (Keyes et al. 1977). To promote the PETT scanner, researchers radically redesigned the scanner and focused on a virtue neglected by the medical imaging community: speed. In 1978, they made the PETT scanner the only device fast enough to image biological processes in the brain in vivo (Ter-Pogossian 1981). They further prepared new questions about cognitive functions, and came up with standardized methods and procedures to use the PETT scanner to answer those questions. In two years, the number of journal articles that included research using the PETT scanner took off – first hundreds and then thousands per year, creating a new field now known as neuroimaging.

While it is true that fierce competition sometimes leads one research program to blot another out, the creation of neuroimaging demonstrates that sometimes fierce competition can change the competitive landscape, modify or even create evaluative criteria, and sometimes completely rewrite the rules of the game. In the case of neuroimaging, competition created an entirely new field with its own rules and virtues.

The creation of neuroimaging further challenges the following philosophical positions. First, no universal or historically stable list of scientific virtues exists. Competition can create new virtues. Second, scientific and pragmatic virtues are deeply intertwined in the history and our evaluation of scientific development. Third, because the rules of the game change constantly, we can only assess local, but not general progress in scientific development.

Keyes Jr, J. W., Orlandea, N., Heetderks, W. J., Leonard, P. F., & Rogers, W. L.

(1977). The Humongotron--a scintillation-camera transaxial tomograph. *Journal of nuclear medicine: official publication, Society of Nuclear Medicine*, 18(4), 381-387.

Kuhn, T. S. (2012). *The Structure of Scientific Revolutions*, 50th anniversary ed. University of Chicago Press.

Lakatos, I. (1970). Falsification and the Methodology of Scientific research Programmes. *Criticism and the Growth of Knowledge*, ed. I. Lakatos and A. Musgrave (London: Cambridge University Press, 1970).

Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. (1996). Niche construction. *The American Naturalist*, 147(4), 641-648.

Ter-Pogossian, M. M. (1981, January). Special characteristics and potential for dynamic function studies with PET. In *Seminars in nuclear medicine* (Vol. 11, No. 1, pp. 13-23). WB Saunders.

Feyerabend's Well-Ordered Science: How an Anarchist Distributes Funds

Jamie Shaw – University of Western Ontario, Canada

To anyone vaguely aware of Feyerabend, the title of this paper seems oxymoronic. For Feyerabend, it is often thought, science is an anarchic practice with no discernible structure. Against this trend, I argue that Feyerabend's pluralism, once suitably modified, provides a plausible account of how to organize science which is superior to contemporary accounts.

Ever since the foundation of the National Science Foundation in 1950, there has been little philosophical analysis of how resources should be distributed amongst diverse and, often, competing research programs. In *Science, Truth, and Democracy*, Kitcher introduces the notion of a 'well-ordered science' where he provides his understanding of how science should be organized. In a nutshell, he posits that democratic deliberations should determine which theories are pursued and how they are prioritized. Since then, others have introduced more fine-grained models that, unwittingly, make use of Kitcher's conception of a well-ordered science (Strevens 2003; Weisberg and Muldoon 2008; Zollman 2010). However, these models conflate the goals of research and the means of attaining those goals. This conflation is the result of assuming that the goals, plus our current scientific knowledge, determines the means of attaining them. For example, if a cure for cancer is a goal of research, we should increase funds for lines of research that seem 'promising' for finding such a cure where the 'promise' comes from our existing knowledge of cancer and its possible treatments (e.g., various subfields of oncology, radiation therapy, etc.). Against this, I argue that Feyerabend was correct in asserting

that we should pursue theories that contradict currently accepted knowledge and appear to have no initial practical value. Therefore, the attainment of the goal (a cure for cancer) also requires funding research that conflicts with current background knowledge (e.g., music therapy) and research that appears to have nothing to do with cancer research by our current lights. In my talk, I will reconstruct the methodological argument Feyerabend provides for this view and show how it supported by the social scientific literature on theory pursuit which shows how solutions to problems came from unexpected sources (Roberts 1989; Foster & Ford 2003; McBirnie 2008).

After this, I go on to show how Feyerabend's pluralism can provide an alternative method of organizing research. Feyerabend's pluralism is essentially the combination of the principle of proliferation, which asserts that we should proliferate theories that contradict existing theories, and the principle of tenacity, which asserts that we can pursue theories despite their theoretical and empirical difficulties indefinitely (Shaw 2017). However, Feyerabend provides no means for balancing these principles and, therefore, his own well-ordered science is incomplete. I argue that this balance can come from C.S. Peirce's 'economics of discovery' which provides limits to the principles of proliferation and tenacity (cf. McKaughan 2008). I conclude by gesturing at recent work on the economics of theory pursuit that provides empirical confirmation of this view (Stephan 2012).

Predicting efficiency of scientific performance in high energy physics

Vlasta Sikimic – University of Belgrade, Serbia

Scientific laboratories in many disciplines have become large and complex. In high energy physics (HEP), laboratories often have hundreds of in-house members, while the number of collaborators working on experiments is measured in thousands. The question of optimization of such laboratories has become of practical importance. While the prevailing approach in social epistemology of science was based on modelling abstract hypothetical scenarios e.g. (Borg et al. 2017, Kitcher 1993, Rosenstock et al. 2016, Zollman 2010), HEP laboratories and their organization have also been analyzed based on actual data, e.g. citation metrics (Martin & Irvine 1984a, 1984b, Perović et al. 2016). The relevant project data are the number of researchers and research teams, the project duration, its citation impact, etc. Benefits of a data-driven approach are the unambiguous interpretation of the results, the predictive power, and the corrective potential when it comes to real-life decisions in science. Still, analyses based on project data can be successfully

applied only under specific terms. For instance, the citation metrics is field-dependent. This becomes obvious already by comparing impact factors of the prominent journals across disciplines. Also, in different scientific fields, specific team size choices can be justified. Finally, it is questionable whether the citation metrics is an informative parameter in a data-driven analysis of a specific field. In HEP, the consensus about the results is relatively quick and stable over long periods of time (decades). The reason for this is the regular inductive behaviour of the field, which postulates the conservation principles as the core ones. Moreover, the Formal Learning Theory approach demonstrates that the stable consensus is a result of a reliable pursuit (Schulte 2000). Thus, the inductive behaviour of the field guarantees the successful and meaningful applicability of data-driven analyses based on citation metrics.

After arguing in favour of a data-driven approach to optimization questions in HEP, I will present an exploratory pilot study based on relevant data concerning the structure and outcomes in HEP. The method used for this efficiency evaluation consists of two stages. In the first stage, data envelopment analysis was conducted on the data from 50 projects run in Fermilab. In the second stage, predictive analysis based on the gradient tree boosting algorithm was applied. This predictive analysis shows how efficient an individual experiment characterized by the relevant parameters will be, and calculates the accuracy of such predictions. This allows us to investigate and analyze trade-offs between input parameters (length of the project, the number of researchers, and the number of teams) with respect to their output (citation counts). The results suggest that projects with lower input values are more efficient. If experiments take shorter periods of time (less than 500 days), the number of researchers plays a smaller role. Moreover, irrespective of the other parameters, researchers should be divided in as few teams as possible. This type of analysis is informative both when it comes to deciding how to structure human resources within a project and when it comes to theoretically analyzing optimal team structures in the field.

Borg AM., Frey D., Šešelja D., & Straßer C. (2017) In S. Benferhat, K. Tabia, and M. Ali, *Advances in Artificial Intelligence: From Theory to Practise: 30th International Conference on Industrial Engineering and Other Applications of Applied Intelligent Systems*, 507–510.

Kitcher, P. (1993). *The advancement of science*. New York: Oxford University Press.

Martin, B. R., & Irvine, J. (1984a). CERN: Past performance and future prospects: I. CERN's position in world high-energy physics. *Research Policy*, 13(4), 183–210.

Martin, B. R., & Irvine, J. (1984b). CERN: Past performance and future prospects: III. CERN and the future of world high-energy physics. *Research Policy*, 13(4), 311–342.

Perović, S., Radovanović, S., Sikimić, V., & Berber, A. (2016). Optimal research team composition: data envelopment analysis of Fermilab experiments. *Scientometrics*, 108(1):83–111.

Rosenstock, S., O'Connor, C., & Bruner J. (2016) In epistemic networks, is less really more? *Philosophy of Science*, 234–252.

Schulte, O. (2000). Inferring Conservation Laws in Particle Physics: A Case Study in the Problem of Induction. *The British Journal for the Philosophy of Science*, 51(4), 771-806.

Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72(1), 17–35.

Can't see the hive for the bees: the importance of biological individuality in behavioural ecology

Jules Smith-Ferguson – University of Sydney, Australia

Madeleine Beekman – University of Sydney, Australia

The question 'what is an individual' does not often come up in studies within the field of behavioural ecology. Generally behavioural ecologists do not think about what makes an individual because they tend to use intuitive working concepts of individuality. Rarely do they explicitly mention how individuality affects their experimental design and interpretation of results. In contrast, the concept of individuality continues to intrigue philosophers of biology. This reaffirms the (often incorrect) stereotype of the scientist doing science, while the philosopher thinks about doing science. A recent review on the philosophical understanding of biological individuality (Pradeu, 2016) highlighted the need for more practical and experimental considerations, in the hope that the philosopher's understanding of biological individuality might benefit from understanding how scientists approach and consider what a biological individual is in their day to day work.

We will try to give such a scientific perspective, by discussing how the biological individuality concept affects experimental considerations in biological fields both within and outside of evolutionary studies (yet with a particular focus on behavioural ecology). In doing so we will touch on a range of different organisms and study questions, with a particular focus on unconventional organisms used in behavioural ecology research. Including unconventional organisms makes it clear why, usually, the concept of biological individuality is intuitive in behavioural ecology. We further explore the reasons

why questions of biological individuality are so often seen through an evolutionary lens. Our ultimate aim is to assist philosophers of biology in their understanding of biological individuality by illustrating the ways in which behavioural ecologists choose their biological individual.

A Synthetic Approach to Studying Scientific Problems

Beckett Sterner – Arizona State University, United States

Scientists widely present and understand their research as aiming to solve specific research problems. Despite its ubiquity, however, the concept of a research problem lacks a synthetic account in contemporary philosophy of science. Historically, Thomas Nickles and Larry Laudan presented important analyses of the nature of problems in the 1970-80s, but these works now appear limited by the theory-centric focus of philosophy of science at the time. In 1973, computer scientist and psychologist Herbert Simon published his famous paper, “The Structure of Ill-Structured Problems,” which spawned an extensive literature on problem structures in organization theory. More recently, philosophers working on scientific integration have made the idea of problems central to their work. Alan Love, for example, introduced the idea of “problem agendas” to help characterize how the field of evolutionary development is generating integration across multiple biological disciplines, and Sandra Mitchell has argued that scientific models can be integrated for the sake of solving specific research problems without leading to general theoretical unification. Work on integration has largely proceeded in isolation, however, from the previous two literatures I mentioned about scientific problems. An important exception is Nancy Nersessian’s adaptation of “problem space” from Simon to characterize the embodied cognitive situations of researchers in laboratories.

This paper aims to synthesize these different bodies of work and show how recent ideas from the integration literature can generalize to have much broader value for characterizing scientific practice. One important obstacle to synthesis is the reliance of Nickles, Laudan, and Simon on formalized scientific theories to characterize important types and features of problems. I argue that research problems are better understood as social institutions that provide crucial common knowledge necessary for researchers to make collective progress toward a solution. I show how Simon’s formal notion of problem structure can be reinterpreted in terms of the performative structuring of problems by researchers. I then discuss how theoretical terms such as problem space and problem agenda can be understood as providing a methodological framework enabling an external observer to articulate and analyze the

meanings scientists ascribe to their activities. I link this framework to grounded theory and the idea of sensitizing concepts in qualitative social science.

In order to show how the idea of structuring research problems generalizes beyond the context of integration, I apply it to an interesting contrast case from the history of systematics. Despite fundamental disagreements and no-holds-barred competition between two groups in systematics, pheneticists and cladists, I show that they nonetheless were able to learn from each other's work in virtue of sharing certain key methodological problems in common. This exchange of knowledge remained possible over time because both groups structured their workflow for classification and phylogenetic inference in similar ways despite their other disagreements. I also show how to use problem structuring to articulate several types of methodological progress that systematists made during this period.

Looking Inside the Black Box: A New Kind of Scientific Visualization

Michael Stuart – London School of Economics and Political Science, United Kingdom

Nancy Nersessian – Harvard University, United States

Computational systems biology is a hybrid field that uses computational methods to address questions about the dynamics of complex biological systems not accessible to traditional biological experimentation. When applied to model the dynamics of complex biological systems, computational systems biology has enjoyed a degree of predictive and explanatory success. However the models themselves are “epistemically opaque,” not just in the usual sense that no human could check all their inferences (Humphreys 2004, Lenhard 2018), but also in the sense that the formal complexity, long length, and idiosyncrasies in coding make it very difficult for others – modellers or biologists – to grasp their content and behavioural profile. This makes these models difficult to understand, interpret, and trust.

In a qualitative study of a computational systems biology laboratory, we found scientists tackling this version of epistemic opacity using visualizations, which is a common technique for increasing understanding (de Regt 2014). We present a kind of scientific visualization they developed, which has not been discussed in the philosophical literature. The visualization is novel in that it is the automatic output of a program designed to generate diagrams of the inner workings of computer models at a slice in time. In other words, the output visualization is a representation of the model, rather than of the biological system being modelled. This kind of visualization successfully resolves the lab's specific problem of opacity, in part because the diagram simplifies and draws

attention to the important parts of the model. But the fact that it is automatically generated adds another layer. Usually such model-diagrams are drawn by hand, and there is no way to verify how well the diagram captures the dynamics of the model. In this case, the diagram is produced by means of a diagram-generating algorithm applied to the model, which can be verified for simple cases. When used in more difficult cases, scientists receive a reliable representation of the computer model's structure and dynamics without the need for a human rendering. We argue that the epistemology of this kind of visualization is best conceived as a powerful form of "non-metaphorical" exemplification (Elgin 2011). Like (Kendall Walton's view of) snapshots, which are photos that represent objects while at the same time providing epistemic access to those objects (Walton 1984), these visualizations offer both informative representations and direct epistemic access. By providing epistemic access, the computer-generated diagrams increase understanding of and trust in the models they portray.

de Regt, H. 2014. "Visualization as a Tool for Understanding." *Perspectives on Science* 22: 377-396.

Elgin, C. Z. 2011. "Making Manifest: The Role of Exemplification in the Sciences and the Arts." *Principia: An International Journal of Epistemology* 15: 399-413.

Humphreys, P. 2004. *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. Oxford: OUP.

Lenhard, J. 2018. "Thought Experiments and Simulation Experiments: Exploring Hypothetical Worlds." Pp. 484-497 in M. Stuart et al. (eds) *The Routledge Companion to Thought Experiments*. London: Routledge.

Walton, K. 1984. "Transparent Pictures: On the Nature of Photographic Realism." *Critical Inquiry* 11: 246-277.

Scientific Pluralism in Practice

David J Stump – University of San Francisco, United States

The philosophical discussion of scientific pluralism has been rather far removed from practical issues. Epistemology and metaphysics dominate the discussion of whether "the ultimate aim of science is to establish a single, complete, and comprehensive account of the natural world" (Kellert, Longino, and Waters 2006, x), or whether science should maintain multiple strategies for describing the world. For example, Dupré and Cartwright make a metaphysical claim that the world itself is plural and not reducible to a single correct interpretation (Dupré 1993; Cartwright 1999), while Ruphy (2016) argues for pluralism epistemologically and says that metaphysical claims are not justified. I would

like to consider the practical consequences of scientific pluralism and begin to sketch ways to proceed from a pluralist point of view. In fact, there are tremendously important practical issues to be dealt with from the standpoint of pluralism. What kinds of research should be pursued and funded? How will decisions on what is pursued and funded be made? These practical questions have yet to be addressed in the literature on scientific pluralism. I will suggest some parameters and strategies for dealing with practical issues from a pluralist point of view. A commitment to pluralism does not mean that every research program should be pursued and funded. A pluralist is quite capable of making judgments about the soundness and viability of proposals. To be sure, one would expect a pluralist to be more open to alternative ideas, but it does not imply that every proposal should be funded equally. This would be a commitment to relativism, which is explicitly denied by many pluralists (Chang 2012, 261). Even assuming that pluralists are in a position to judge individual proposals as good or bad, there is a further question of whether or not they can consistently make judgments about entire areas of research, rather than judging proposals on their individual merit. I would argue that a pluralist can dismiss a research tradition, just as they can judge individual proposals. Thus, areas of research that are frequently labeled as pseudoscience might be rejected, even from a pluralist viewpoint. What then is a commitment to pluralism? Does it mean rejecting reductionist schemes out of hand? Or can a pluralist even accept that reductions occasionally work in limited areas? This talk explores how to mediate between the extremes of “anything goes” and the “winner take all” in scientific decision making.

Motivated Thinking About Behavioral Genetics: What It Is And What To Do About It

Kathryn Tabb – Columbia University, United States

Matthew Lebowitz – Columbia University, United States

Popular television shows, novels, and science writing about forensics and the law often rely on a widely-held intuition: that when behaviors are discovered to be caused by genetic dispositions, they become less blameworthy. And in real-life contexts like the courtroom, genetic information is increasingly brought to bear in order to mitigate punishments for antisocial or violent crimes. Yet, psychologists have not been able to produce this intuition in the laboratory setting as robustly as one might expect – indeed, there is very little empirical work showing any such effect at all, with most studies showing genetic information to have no effect on moral judgments.

In this talk we present some novel and robust empirical data suggesting why this might be the case. Using evidence from several studies, we show that people seem unwilling or unable to accepting genetic explanations for antisocial behavior, while comparable explanations for prosocial behavior are accepted without equivalent resistance. In other words, before experimental subjects even get to the step of integrating genetic information about bad actions into their assessments of moral responsibility, they have already discounted its worth.

Our first aim is to determine what might cause this asymmetry. We consider what we view as the two most likely reasons why people might respond differently to genetic information about prosocial and antisocial behavior in their moral reasoning. On the one hand, people might reject some genetic explanations because of genetic determinism, the intuition that genes determine behavior, and thus mitigate free will. Such beliefs, we show, might lead people to reject information that would inhibit their ability to punish others. This explanation is in line with research showing that people are less likely to see behavior as determined in situations where they have a desire to hold others accountable. On the other hand, people might reject genetic explanations of undesirable behaviors because of genetic essentialism, the intuition that genomes constitute who we really are. Research has shown that people avoid judgments that condemn the essential natures of others, preferring to see bad behavior as inconsistent with who others are deep down. This research would suggest that insofar as genes are seen as defining our deepest selves, we would be uncomfortable with the idea of genes “for” evil dispositions.

Our second aim is to consider the methodological and philosophical repercussions for our findings. If experimental-philosophical or psychological studies seek to assess how moral responsibility attributions are made on the basis of genetic information, they will need to take the asymmetry we have brought to light into account. The asymmetry also needs to be considered in debates over whether behavioral genetics has a place in legal arguments about competency and culpability. On the theoretical side, insofar as our argument suggests that people’s moral judgments take primacy over scientific information in assessments of moral responsibility, there is a normative question of whether or not they should. We end by considering whether the way people seem to integrate scientific and moral information when attributing blame is appropriate, or worrisome.

Robustness Analysis in Network Neuroscience

Morgan Thompson – University of Pittsburgh, United States

By applying the mathematical modeling tools of graph theory and network science to neuroscience data, network neuroscientists have raised new questions and methodological issues in neuroscience. Network neuroscientists aim to model neural systems as networks of nodes (or parts) and the connections between them as well as to characterize their organization. Network neuroscientists address two fundamental questions: What topological properties characterize the organization of neural systems? And how do neural systems come to be organized in those ways? Robustness analysis is particularly helpful in aiding researchers in answering these questions. Robustness analysis is the process of searching for a robustness relationship between some property or prediction and a set of models. In this talk, I distinguish between two types of robustness analysis using case studies from network neuroscience to illustrate the conditions under which each type should be deployed, can fail, or is misused.

Robustness analysis can be used to identify whether some property is robust over a pre-established set of models with varied assumptions. This process helps with the first research question in network neuroscience. Before the question can be answered, neuroscientists must determine how to appropriately model the neural system as a network (Bassett and Sporns 2017). The methodological choices in constructing the network model from neuroimaging data might influence the resulting topological properties (Zalesky et al 2010). I illustrate cases in which network neuroscientists use robustness analysis on a predefined set of models with the methodological assumptions under debate (e.g., different atlas-based or defining nodes) and determine whether the models identify the same topological properties. In cases of failure to identify the same topological property for all the models, researchers use the opportunity to reflect on the methodologies employed. I demonstrate that this type of robustness analysis is well poised for methodological reflection and need not be understood in the context of confirmation, which has been the focus of previous work on robustness analysis (e.g., Orzack and Sober 1993; Lloyd 2015; Parker 2011; Houkes and Vaesen 2012; Schupbach 2015).

The other type of robustness analysis takes properties of interest and determines which set of models have that property. While this latter type of robustness analysis has been under-described in the philosophical literature on robustness analysis (e.g., Weisberg 2013), it is most appropriate for addressing the second research question in network neuroscience. Neuroscientists use

generative principles (e.g. wiring rules; see also Green 2013) that are plausible based on existing theory and data to create synthetic models of the target system. These models can then be examined for the properties of interest known from descriptive models of the target system; when a set of principles cannot generate models with these properties of interest, the principles can be eliminated as possible developmental principles for the target system. A well-accepted descriptive model (e.g., Ankeny 2006) of the target system such as the *C. elegans* wiring diagram (White et al. 1986; Varshney et al. 2011) is necessary for this type of robustness analysis to play this role.

Ankeny, Rachel. 2007. "Wormy Logic: Model Organisms As Case-Based Reasoning". In Creager, A.N.H., Lunbeck, E., and Wise M.N. *Science without Laws: Model Systems, Cases, Explanatory Narratives*. Durham: Duke University Press. P. 46-58.

Green, Sara. 2013. "Tracing Organizing Principles: Learning from the History of Systems Biology." *Hist. Phil. Life Sci.*, 35: 553–76.

Houkes, Wybo, and Krist Vaesen. 2012. "Robust! Handle with Care." *Philosophy of Science* 79 (July): 1–20.

Lloyd, Elizabeth. 2015. "Model Robustness as a Confirmatory Virtue: The Case of Climate Science." *Studies in History and Philosophy of Science* 49: 58–68.

Orzack, Steven Hecht, and Elliott Sober. 1993. "A Critical Assessment of Levins's *The Strategy of Model Building in Population Biology* (1966)." *The Quarterly Review of Biology* 68 (4): 533–46.

Parker, Wendy. 2011. "When Climate Models Agree: The Significance of Robust Model Predictions." *Philosophy of Science* 78 (4): 579–600.

Schupbach, Jonah N. 2015. "Robustness, Diversity of Evidence, and Probabilistic Independence." In *Recent Developments in the Philosophy of Science: EPSA13 Helsinki*, edited by Votsis Maki and Schurz Ruphy, 305–16.

Varshney, Lav R, Beth L Chen, Eric Paniagua, David H Hall, and Dmitri B Chklovskii. 2011. "Structural Properties of the *Caenorhabditis Elegans* Neuronal Network." *PLoS Computational Biology* 7 (2).
doi:10.1371/journal.pcbi.1001066.

Weisberg, Michael. 2013. *Similarity and Simulation: Using Models to Understand the World*. Oxford University Press.

White, J.G., E. Southgate, J.N. Thomson, and S. Brenner. 1986. "The Structure of the Nervous System of the Nematode *Caenorhabditis Elegans*." *Philos.Trans.R.Soc Lond B Biol.Sci.* 314: 1–340.

Zalesky, Andrew, Alex Fornito, Ian H. Harding, Luca Cocchi, Murat Yucel, Christos Pantelis, and Edward T. Bullmore. 2010. "Whole-Brain Anatomical Networks: Does the Choice of Nodes Matter?" *NeuroImage* 50 (3). Elsevier Inc.: 970–83. doi:10.1016/j.neuroimage.2009.12.027.

Doing science for practice: lessons from medical imaging development

Sophie van Baalen – Universiteit Twente, Netherlands

In traditional philosophy of science, application of science to practice is considered to be a relatively straightforward process that primarily involves deductive reasoning from general laws. This view was challenged by philosophers of science such as Cartwright (1999) and Toulmin (2001), according to whom it is a myth that scientific products (e.g. theories, concepts, principles, laws) are universal and detached from the experimental and theoretical context.

The application of science outside of the initial domain requires contextualization by interpretation or extrapolation. Furthermore, the solution of real-world problems usually requires input from multiple scientific disciplines. At the same time, producing epistemic products intended for a specific practice places requirements on these products to make them intelligible and relevant to the situation (e.g. Boon, 2017 and De Regt, 2015). These aspects make the application of science challenging for scientist as well as professionals using scientific products in practice. The aim of this paper is to illustrate these issues in the context of two case studies concerning the development and introduction of new magnetic resonance imaging (MRI) tools for clinical practice, one in which I was involved as clinical researcher.

The development of an imaging tool for clinical practice requires combining multiple disciplinary perspectives. Disciplinary perspectives involve a coherent, interdependent set of epistemic elements that are specific to the discipline and guide epistemic activities such as formulating research questions, discerning specific types of phenomena, formulating hypotheses and conceptualizing processes, mechanisms or objects that cannot be directly observed. An important challenge of this type of work is combining the elements in such a way that the product is logically coherent as well as consistent with all perspectives that are involved.

In the context of medical imaging development close collaboration between experts with a different disciplinary background, i.e. medical and technological, is needed. For a new tool to provide clinically relevant information (that is, information that may lead to a change in clinical outcome) it is necessary to establish what features are visible in the images, and which are relevant for clinical practice. The ability to recognize relevant features co-evolves with the development of the technology itself and it is therefore crucial to have a direct feedback cycle between clinicians and image producers such as radiographers and researchers. This allows making small adjustments to (“tinkering with”) the

technology and establishing a way of looking at the images in an ongoing exchange between clinician and technology developer. To make sense of a new imaging tool clinicians have to gain experience with it in a wide range of clinical situations, in order to relate the technical outcome to the clinical outcome, and hook it onto something that is familiar, such as another – established – imaging modality or other clinical measurement.

Finally, reflecting on these case studies I will comment on methodology for studying the application of science in professional practice.

Toulmin, S. (2001). *Return to reason*

Cartwright, N. (1999). *The dappled world*

Boon, M. (2017). *Philosophy of science in practice: A proposal for epistemological constructivism*

De Regt, H. W. (2015). *Scientific understanding: truth or dare*

Rethinking the explanatory power of dynamical models in cognitive science

Dingmar van Eck – Ghent University, Belgium

In this paper I offer an interventionist perspective on the explanatory structure and explanatory power of (some) dynamical models in cognitive science: I argue that some ‘pure’ dynamical models – ones that do not refer to mechanisms at all – in cognitive science are ‘contextualized causal models’ and that this explanatory structure gives such models genuine explanatory power. I contrast this view with several other perspectives on the explanatory power of ‘pure’ dynamical models. One of the main results is that dynamical models need not refer to underlying mechanisms in order to be explanatory, as some authors maintain (Kaplan and Bechtel 2011). I defend and illustrate this position in terms of dynamical models of the A-not-B error in developmental psychology as elaborated by Thelen, Smith, and collaborators (2001), and dynamical models of unintentional interpersonal coordination developed by Richardson, Schmidt, and collaborators (2007, 2008).

Contextualized causal model explanations are comprised of causal claims that cite core causal factors that make a difference to explananda phenomena and descriptions of (internal and external) constraints and their relations with core causal factors, which specify when core causal factors are (and when they are not) difference makers for explananda phenomena. These contextual dependencies between constraints and core causal factors are counterfactual dependencies, which give such explanations genuine explanatory power (despite the fact that they do not specify mechanisms).

In the case of the A-not-B error this works as follows. This error concerns the phenomenon, seen in 7-12 months old infants, where infants continue to reach to a location where they have uncovered a toy even when they see in subsequent trials the toy being hidden at a new location. Motor memory is key to the explanation of Thelen et al. (2001, p. 9): “infants make perseverative location errors because the motor memory of one reach persists and influences subsequent reaches”. This factor is represented in a dynamical model as coupled with a number of internal and external constraints. Specifications of these couplings are key to the model’s explanatory power: phrased in (broadly) Woodwardian (2003) terms, the model of Thelen et al. (2001) answers what-if-things-had-been-different questions by spelling out the mutual interdependencies between the core causal factor of motor memory of previous reaches and internal constraints such as looking, motoric planning, and reaching, and external ones like the relative ambiguity of the task input and the delay between looking and reaching. The model characterizes under which contextual conditions motor memory of previous reaches makes a difference to perseverative reaching, and when it does not. For instance, how constraints like body posture and the salience of the hiding locations affect whether or not motor memory of previous reaches is a difference maker for perseverative reaching. These dependencies tell us how changes to the values of the constraints in these dependencies result in changes in the value of the target explanandum phenomenon, i.e., they tell us under which value changes, motor memory is and when it isn’t a difference maker for perseverative reaching.

Alongside this analysis I offer principled reasons why contextualized causal models are distinct from mechanistic models, thus countering the possible objection that the former ones are mechanistic (and therefore explanatory!) after all.

Kaplan, D.M. & Bechtel, W. (2011). Dynamical models: an alternative or complement to mechanistic explanations? *Topics in Cognitive Science* 3: 438-444

Richardson, M.J., Marsh, K.L., Isenhower, R.W., Goodman, J.R.L., & Schmidt, R.C. (2007). Rocking together: Dynamics of intentional and unintentional interpersonal coordination. *Human Movement Science* (26): 867-891.

Schmidt, R.C., & Richardson, M.J. (2008). Dynamics of interpersonal coordination. In *Coordination: neural, behavioral, and social dynamics* (Eds. Fuchs, A. and Jirsa, V.K.). Springer-Verlag, Berlin Heidelberg: 281-308

Thelen, E., Schöner, G., Scheier, C., & Smith, S.B. (2001). The dynamics of embodiment: a field theory of infant perseverative reaching. *Behavioral and Brain Sciences* 24: 1-34

Woodward J. (2003), *Making Things Happen. A Theory of Causal Explanation*. New York: Oxford University Press.

A role for the history and philosophy of science in the promotion of scientific literacy

Holly VandeWall – Boston College, United States

In a democratic system non-experts should have a voice in research and innovation policy, as well as in those policy issues to which scientific and technological expertise are relevant – like climate change, GMOs and emergent technologies. The inclusion of non-expert voices in the debate is both a requirement for truly democratic process and an important counter to the kinds of jargon and group-think that can limit the perspective of more exclusively expert discussions.

For non-experts to participate in a productive way does require a certain degree of scientific literacy. Yet in our present place of intensive specialization, access to understanding any one subfield or subdiscipline in the sciences requires years of study. Moreover, the relevant sort of literacy involves not simply familiarity with factual information, but some perspective on the goals, methods and practices that constitute knowledge formation in the scientific disciplines.

We want to have a conversation about how those who teach history and philosophy of science related courses might engage course material in ways that more effectively promote science literacy of a kind that provides non-experts in the sciences an entry point to the necessary deliberation of how we fund, promote, limit and oversee scientific research in our society.

Drawing on our own experience teaching non-majors using primary sources from the history and philosophy of science from Plato to NATO, we will offer some specific suggestions for how studying the history of science can contribute to the socially relevant goal of fostering science literacy in practical ways, without compromising responsible attention to historical context and scientific ideas.

In particular, we have found that coursework that familiarizes students with the how practices of knowledge formation in the sciences have developed over time has helped our students to:

1. Recognize that the methods of science are themselves developed through trial and error, and change over time
2. Understand that different disciplines of science require different approaches and techniques, and will result in different levels of predictive uncertainty and different standards for what is considered a successful hypothesis

3. Consider examples of scientific debate and processes through which those debates are resolved with the advantage of historical perspective
4. Trace some of the unintended effects of the sciences on society and to identify where the social and cultural values of the scientists themselves played a role in their deliberations – and whether or not that had a negative epistemic effect.

Who knows which way the evidence is going?

Veronica Vieland – Ohio State University, United States

Various methods exist for measuring the strength of statistical evidence (MaxLR, Bayes factor, p-value, inter alia). They can disagree with one another regarding which of two hypotheses is favored by the data, or whether the evidence is increasing or decreasing as data accrue. Thus they cannot all be measuring the same thing. Given the nearly universal interpretation of statistical results in evidential terms, this seems problematic. Yet there is little agreement on how to go about remedying it. How would we establish which if any of these candidate evidence measures is actually measuring the evidence?

Measurement requires concatenation operations, i.e., rules for obtaining amalgamated measurements from a set of individual measurements. Therefore one way to compare alternative evidence measures is in terms of their amalgamation behavior. If we can agree on how the evidence is behaving as data accrue, then we can use this to assess whether the evidence measure tracks properly with the evidence. For example, consider the strength of evidence that a defendant was present at the scene of a crime based on a blood type match (“weak” evidence that the defendant was there) and a DNA match (“stronger” evidence). We would all agree that the evidence in aggregate is at least as strong as the DNA evidence alone, that is, the blood type match does not weaken the evidence provided by the DNA match. This has been proposed elsewhere as an evidence amalgamation principle (EAP): If two data sets $D1$ and $D2$ each support the same hypothesis (over some specified alternative), then the amalgamated evidence $Ev(D1, D2) > \text{maximum}[Ev(D1), Ev(D2)]$.

The EAP seems simple enough, but the legal example involves only rank-ordering, not quantification, of evidence strength, and it is unclear how this relates to statistical reasoning. Indeed, most familiar statistical evidence measures routinely violate the EAP, returning $Ev(D1) < Ev(D1, D2) < Ev(D2)$ under a broad range of circumstances. Still, there is something compelling about the EAP itself. So which is correct when it comes to evidence amalgamation, the EAP or standard statistical methodology?

Here we distinguish between two modes of evidence amalgamation: accumulative and agglomerative. Accumulation of evidence involves the use of new observations to revise previous assessments of evidence strength; agglomeration leaves individual (per-assay or per-study) assessments of evidence in tact when considering the total evidence on all available information. It could be argued that the EAP applies in the context of agglomerative, but not necessarily accumulative, evidence amalgamation. Virtually all statistical theory presupposes accumulation as the standard mode, explaining our difficulty in relating the EAP to statistical settings. But we will illustrate that common approaches to evidence amalgamation, including meta-analysis, fail to exhibit coherent behavior in light of the accumulative-agglomerative distinction. Indeed, even in accumulative contexts standard methods can return clearly erroneous results, e.g., $Ev(D1, D2) < \text{minimum}[Ev(D1), Ev(D2)]$. Thus even after refining the conditions of applicability of the EAP, we still lack a measure of statistical evidence that satisfies minimal criteria for empirical adequacy.

Inventing Units of Measurement: Lessons from Newton's Optics

Kirsten Walsh – University of Nottingham, United Kingdom

Anyone who has studied a small amount of classical mechanics will be familiar with the newton: the metric unit of force named after Sir Isaac Newton. It's a useful unit if you want to go rock climbing or fly a fighter jet. Fewer people have heard of the interval of fits—invented by Newton to explain why a body reflects light of one colour rather than another colour. This unit of measurement allowed Newton to offer both a quantitative theory of coloured bodies (his 'theory of fits') and an instrument for measuring extremely small things (in the order of 1/100,000th of an inch). The theory of fits is not recognised as one of Newton's greatest achievements—partly because, abstracted and formalised in book 2 of the *Opticks*, it was nearly incomprehensible. I argue, however, that the process by which Newton invented the interval of fits and eventually arrived at his theory of fits is revelatory of the role of metaphysics in Newton's natural philosophy.

In his investigations of interference phenomena (now known as 'Newton's rings'), Newton postulated hypothetical æthereal vibrations, produced when a ray of light passes through a surface (like the ripples produced when a stone is thrown into a pond), in order to explain the periodic reflection and transmission of light. The ray of light causes the æther to expand and contract—in other words, to 'pulse'—in a regular pattern. By establishing the connection between the dimensions of thin films and the rays that are

reflected or transmitted, Newton was able to measure the length of a pulse. Operationalising the concept of a pulse—i.e. defining the concept through the operations by which it was measured—gave Newton a unit of measurement and, eventually, a way of formalising and abstracting the explanation.

I draw two lessons from this case study. The first concerns Newton's use of hypotheses. Newton was explicitly and (in)famously anti-hypotheses. This can be seen, for example, in the oft-quoted passage from the General Scholium: *Hypotheses non fingo*. And yet, while hypotheses were certainly not the goal of this investigation—they were not an end in themselves—they nonetheless played a vital supporting role in the development of Newton's theory of fits.

The second lesson concerns how we should understand Newton's metaphysical theses. They are not as concrete as they first appear. Newton's hypotheses about the natures of things were based upon the capacities of certain kinds of entities and processes to support properties he had experimentally established. And so, we should understand Newton's metaphysical posits functionally: to be a corpuscle, on Newton's account, is to play a certain kind of role. Moreover, unless the hypothesis could be put to some particular use—e.g. suggesting an experiment or offering a possibility proof—he was not interested in developing it further.

Registration Pluralism and the Cartographic Approach to Inter-Individual Brain Differences

Zina Ward – University of Pittsburgh, United States

It is now well known that the overall size of different brains (Rushton & Ankney 1996), their gyral and sulcal geography (Ono et al. 1990), their connectivity patterns (Mueller et al. 2013), the size of their cytoarchitecturally-defined regions (Scheperjans et al. 2008), and the location of those regions relative to macroanatomical landmarks (Amunts et al. 1999) are all highly variable between people. This variation constitutes a significant obstacle to neuroscientific research, which often requires identification of the same part or location across different brains (Brett et al. 2002, Devlin & Poldrack 2007). The neuroscientific community has largely adopted a “cartographic approach” in response to this challenge: cross-brain comparison and aggregation are achieved by placing brain data from multiple subjects into a common reference frame or onto a single template using a registration method (Toga et al. 2006). This mapping of individual data into a shared space allows the “same” place in multiple subjects' brains to be identified and group-level statistics to be run. In this paper, I characterize the cartographic approach and argue that it must be applied in a way that is sensitive to the research question at hand. I begin by

providing a non-technical introduction to the crucial components of the cartographic approach, including stereotactic spaces, templates, atlases, and registration algorithms. I describe briefly how the cartographic approach has changed over the last century: from the visual inspection of paper atlases, constructed from post mortem examination of stained sections, in the early twentieth century; through the invention of stereotactic spaces and early landmark-based alignment methods, which were designed for neurosurgery but coopted in the 1980s for neuroimaging research; to the construction of digital brain atlases and the proliferation of automated and semi-automated registration methods, which continues today (Toga & Mazziotta 2002; Toga et al. 2006; Evans et al. 2012).

On one interpretation of this history, researchers have gotten ever closer to the ideal registration method as techniques have incorporated more types of information and become more mathematically complex. Implicit in this interpretation is the assumption that there exists a single best way to register brains to a template. In the remainder of the paper, I argue against this view. I defend what I call “registration pluralism”: the thesis that there is more than one appropriate way to register a brain to a template. The best registration method in any particular case depends on the phenomenon of interest. This position is an analog of pluralism about atlases, a far more visible and popular position, which states that there are multiple, equally valid ways of dividing up the brain and labeling its parts (Bohland et al. 2009; Amunts et al. 2014). Registration pluralism is in the same spirit. Different registration methods (and not just different partitioning schemes) ought to be applied in different scientific contexts.

After arguing for this view by appeal to the conditions under which registration is judged to succeed, I suggest that registration pluralism has several underappreciated methodological consequences. First, researchers should choose registration methods in a more context-sensitive manner; second, neuroscientists ought to approach both the construction and use of multi-modal atlases with great caution; and third, functional registration, which uses functional information to align brains to a template, has important limitations.

Referential norms and practices in biological taxonomy

Joeri Witteveen – Utrecht University, Netherlands

The norms and practices biological taxonomists follow to fix the reference of taxonomic names have changed over time and continue to differ between domains of biological taxonomy. For example, virologists and zoologists use different standards for establishing what makes something a “valid” taxonomic

name, and for “tracking” which grouping of interest a name refers to as classifications change. In this paper, I examine how these differences in referential practices have come about, before analyzing the epistemic and (meta-)normative implications of the parallel use of different systems of reference.

I begin by arguing that systems of reference in taxonomy are neither mere conventions nor apply one or the other theory of reference because its purported truth. Using a brief comparative analysis of the origination of two different codes of taxonomic nomenclature that are currently in use, I show how norms of reference have been “negotiated” by communities of biologists, by considering the availability of technologies (e.g. for the preservation and of specimens), the realism of classificatory divisions in nature, the size of the communities producing and using scientific names, and the (expected) stability of the current classificatory divisions in the taxonomic domain at issue. I show how these factors, among others, have played into decisions about which referential practices ought to be adopted and institutionalized within a taxonomic community.

In the second part of the paper, I analyze some key implications of the divergent referential practices that have become adopted in different domains of biological taxonomy. Through a comparison of a code of nomenclature for viruses (ICTV) with the codes of nomenclature for plants, fungi and algae (ICN) and zoology (ICZN), I document when and how the scientific practice of forming taxonomic hypothesis becomes constrained by nomenclatural practices. I argue that, though in principle avoidable, this constraining of scientific practice by nomenclature procedures may be unavoidable and have long term benefits, but only under particular conditions – whether these are currently met is should be a matter of concern for practicing taxonomists.

Retraction in Science

K. Brad Wray – Aarhus University, Denmark

Line Andersen – Aarhus University, Denmark

Publication plays a pivotal role in the growth and dissemination of scientific knowledge. But the growth of knowledge is neither strictly linear nor unidirectional. Mistakes are made. Retraction is one means by which the scientific record is corrected. I examine the retraction practices and prevalence in Science. This initial study focuses on the retractions in one single year of published issues of Science.

I report on data on:

the number of scientists who authored the retracted articles; the number of scientists who signed the retraction; how often (i) all authors, (ii) only some authors, and (iii) the editors of Science retracted the article (see Wray 2018); the time between publication and retraction; whether the articles continued to be cited after it was retracted (a previous study by Budd et al. 1998 suggests retracted articles were cited for their reported findings even after they were retracted); and the cause of the retraction, for example, (i) non-replicable results, (ii) manipulated data or images, or (iii) lost data or records, (iv) etc.

My aim is to identify patterns in the retractions, thus revealing a structure to this aspect of the growth of scientific knowledge. This study will aid us in understanding the publication process, and the the role that corrections play in the growth of science. Retraction is just one form of correction in the scientific literature, but a deeper and systematic understanding of it will provide a fully picture of the role of publication in science in advancing scientific knowledge.

Budd, J. M., M. Sievert, and T. R. Schultz. 1998. "Phenomena of Retraction: Reasons for Retraction and Citations to the Publications," *Journal of the American Medical Association*, 280: 3, 296-297.

Cicchetti, D. V. 1991. "The Reliability of Peer Review for Manuscript and Grant Submissions: A Cross-disciplinary Investigation," *Behavioral and Brain Sciences*, 14: 119-186.

Elia, N., E. Wager, M. R. Tramèr. 2014. "Fate of Articles that Warranted Retraction Due to Ethical Concerns: A Descriptive Cross-Sectional Study," *PLOS ONE*, 9: 1. (January 2014), 1-8. e85846.

Wray, K. B. 2018. "The Impact of Collaboration on the Epistemic Cultures of Science," in *Scientific Collaboration and Collective Knowledge*, edited by T. Boyer-Kassem, C. Mayo-Wilson, and M. Weisberg. Oxford: Oxford University Press, pages 117-134.

Wray, K. B. 2006. "Scientific Authorship in the Age of Collaborative Research," *Studies in History and Philosophy of Science*, Vol. 37: 3, pages 505-514.

Wray, K. B. 2002. "The Epistemic Significance of Collaborative Research," *Philosophy of Science*, Vol. 69: 1, pages 150-168.

Making Data Patterns and Interpreting Data in Neuroimaging Research

Jessey Wright – Stanford University, United States

Techniques for analyzing, handling, and sharing data have been central to progress in neuroimaging over the last decade, and to critiques of the technology. While on one hand new analysis technologies have often

motivated theoretical advances, on the other philosophers have issued a variety of arguments against common interpretations of the data. These arguments insist that popular uses of neuroimaging data are undermined by its indirect nature — which reflects blood flow, not neural activity — and the assumptions implicit in the complex data manipulations used to interpret it (Mole and Klein 2010; Aktunc 2014). However, the contribution that data analysis techniques make to inferences in neuroimaging may be more nuanced than these criticisms allow. While critics often treat data analysis techniques as tools for error correction or statistical inference, data manipulations appear at almost all stages of the experimental process, not just in the final moments of data interpretation (Roskies 2010; Israel-Jost 2016; Wright 2017). Here, I contribute to the growing body of work that examines evidence in neuroimaging from a variety of standpoints in order to better understand how the way images are produced and data is manipulated influences the claims and theories the results are treated as evidence about. Drawing upon my experience collaborating with neuroscientists and my current situation as a resident in a neuroscience lab, I propose that data manipulations are primarily used to isolate interpretable patterns in the data. I then argue that decisions about which analysis techniques to use and how to use them are informed by judgements of the interpretability of the patterns they produce.

To develop this view I examine the uptake of multi-voxel pattern analysis (MVPA) techniques, which has led to the use of neuroimaging data to investigate the information represented in brain activity (Tong and Pratte 2012). I argue that one catalyst for this shift is that MVPA techniques became regarded as able to isolate patterns that can be interpreted by claims about the information represented in measured brain activity. To illustrate this I examine a recent dispute over meta-analysis results (Lieberman and Eisenberger 2015; Yarkoni 2015). The dispute is between users of the NeuroSynth database — an automatically curated and analyzed collection of published neuroimaging results — who argue that the data provides strong evidence for a claim, and the database developer who argues their claims cannot possibly be supported by the data. I illustrate how the decision to use NeuroSynth is motivated by a judgement that the resulting data patterns could provide supporting evidence for the claim in question, and that this potential informs the decisions investigators make in the application of the analysis methods and subsequent interpretation of their results. This analysis shows that how data analysis techniques are understood, discussed, and promoted by the community has a significant influence on the evidential value ascribed to data. An influence that can precede the actual manipulation and interpretation of the data itself.

- Aktunc, E. M. [2014]: 'Severe Tests in Neuroimaging: What We Can Learn and How We Can Learn It', *Philosophy of Science*, 81, pp. 961-73.
- Anderson, M. [2015]: 'Mining the brain for a new taxonomy of the mind', *Philosophy Compass*, 10, pp. 68-77.
- Israel-Jost, V. [2016]: 'Computer Image Processing: An Epistemological Aid in Scientific Investigation', *Perspectives on Science*, 6, pp. 669-695.
- Lieberman, M. D., and Eisenberger, N. I. [2015]: 'The dorsal anterior cingulate cortex is selective for pain: Results from large-scale reverse inference', *PNAS*, 112(49), pp. 15250-15255.
- Mole, C. and Klein, C. [2010]: 'Confirmation, Refutation, and the Evidence of fMRI', In Stephen Hanson & Martin Bunzl (eds.), *Foundational Issues in Human Brain Mapping*. Cambridge: MIT Press. pp. 99-112.
- Roskies, A. [2010]: 'Saving Subtraction: A reply to Van Orden and Paap', *British Journal for the Philosophy of Science*, 61, pp. 635-65.
- Tong, F., and Pratte, M. S. [2012]: 'Decoding Patterns of Human Brain Activity', *Annual Review of Psychology*, 63, pp. 483-509.
- Wright, J. W. [2017]: 'The Analysis of Data and the Evidential Scope of Neuroimaging Results', *British Journal for the Philosophy of Science*.
- Yarkoni, T. [2015]: 'Still not selective: comment on comment on comment on Lieberman and Eisenberger (2015)'. Available at: <http://www.talyarkoni.org/blog/2015/12/14/still-not-selective-comment-on-comment-on-comment-on-lieberman-eisenberger-2015/>

Modes of Experimental Interventions in Molecular Biology: A Case Study of the β -Galactosidase Synthesis

Hsiao-Fan Yeh – National Chung Cheng University, Taiwan

Ruey-Lin Chen – National Chung Cheng University, Taiwan

This paper explores modes of experimental interventions in molecular biology. We argue for the following three points: (i) Experimental scientists in molecular biology may use intervention as an essential means to both test hypotheses and to explore (or discover) novel phenomena. (ii) We can distinguish different modes of experimental interventions according to the two standards: the interventional direction and the interventional effect. (iii) A series of related experiments may be conducted to both test and explore by using multiple modes of interventions.

Our argument begins with a brief characterization of Craver and Darden's taxonomy of experiments, because the taxonomy they have made implies different modes of intervention (Carver and Darden 2013). We propose to

extract two interventional directions and two interventional effects as the classificatory grounds from their taxonomy of experiments. The vertical or inter-level direction means that an intervention is performed between different levels of organization and the horizontal or inter-stage direction means that an intervention is performed between different stages of a mechanism. Interventions may produce an excitatory or an inhibitory effect. As a consequence, we can classify modes of intervention according to different interventional directions and effects.

We will do a case study of the PaJaMa experiment (Pardee, Jacob and Monod 1959) to illustrate the three points. The PaJaMa experiment is really a series of experiments on the synthesis of β -galactosidase. Scientists used the series of experiments to both test the endogenous inducer hypothesis and explore a novel property (the inducibility in the case of “double prevention”, i.e., the inhibition of the repressor). Their manipulation of the wild type gene I+ that encodes the repressor protein is a vertical excitatory intervention and the manipulation of the mutant genes I- that does not encode the repressor an inhibitory intervention. A horizontal intervention was taken to produce excitatory effects by adding external inducers, and the other horizontal inhibitory intervention to terminate the synthesis by injecting the wild type gene into female’s cells. Using the multiple modes of interventions together in the series of experiments led to the discovery of the synthesis of β -galactosidase.

Craver, Carl F. and Lindley Darden (2013). *In Search of Mechanisms: Discoveries across the Life Sciences*. Chicago: The University of Chicago Press.

Pardee, Arthur B., François Jacob and Jacques Monod (1959). “The Genetic Control and Cytoplasmic Expression of ‘Inducibility’ in the Synthesis of β -galactosidase by *E. Coli*.” In Andre Lwoff and Agnes Ullmann (eds.) *Selected Papers in Molecular Biology by Jacques Monod*. Academic Press Inc.

Pardee, Arthur (1979). “The PaJaMa Experiment.” In Lwoff, Andre and Agnes Ullmann (eds.) *Origins of Molecular Biology: A Tribute to Jacques Monod*, pp. 109-116. New York: Academic Press.

Schaffner, Kenneth (1974). *Logical of Discovery and Justification in Regulatory Genetics*. *Studies in History and Philosophy of Science* 4:349-385.

Similarity in Practice

Martin Zach – Charles University, Czech Republic

Similarity has long been viewed as the account of scientific representation (e.g. Giere 1988). According to a popular version of the similarity account, scientists utilize similarity relations (in certain aspects to certain degrees) between models and their target systems for representational purposes. However, the account has been challenged. For instance, it has been argued that similarity has the wrong logical properties to be an account of representation, or that similarity is neither necessary nor sufficient for representation (e.g. Suárez 2003; Frigg and Nguyen 2017). Others have claimed that due to the metaphysical nature of non-material models, i.e. they do not instantiate the spatio-temporal properties of their target systems, similarity is a non-starter because one cannot compare a non-instantiated property with one that is instantiated (e.g. Hughes 1997; Thomson-Jones 2010; Odenbaugh 2015). Furthermore, if the target does not exist then there can be no relation of similarity between the model and its target (Toon 2012). Some have pushed back against the objections, suggesting that the critics might have misconstrued the similarity account in some ways (e.g. Godfrey-Smith 2009; Chakravartty 2010; Mäki 2011; Weisberg 2013).

I suggest to consider a different approach, one that consists of two steps. First, I give up on the quest of reinstating the similarity account to its former status as the account of representation, which allows me to escape most of the above objections. Instead, I adopt the notion of representational style (Frigg and Nguyen 2017) where similarity is no longer construed as grounding the notion of representation. In my view, representation is established by an act of stipulation (Callender and Cohen 2006). However, such an act cannot be an arbitrary stipulation, as number of authors have shown (e.g. Frigg and Nguyen 2017). Rather, I introduce the notion of pragmatically and cognitively constrained stipulation (PCCS), building up on insights of Bolinska (2013) and Knuuttila (2017), among others. Pragmatic constraints come from different research goals. Turning a vehicle into a representation is dependent on what the particular aim is (some vehicles are thus more useful than others, some not at all). Cognitive constraints concern our cognitive make-up. Even 'non-material' models often, if not always, have various material dimensions that allow us to have cognitive access to their targets to different degrees.

Second, I argue that the critics of similarity have, perhaps, gone too far with their objections, and as a result, they have underestimated the actual value of similarity judgements not only in scientific practice, but also in ordinary human cognition (e.g. in concept formation). We see the similarity judgments being

employed in various forms; from similarity of properties to similarity of patterns, and from similarity of mechanisms to similarity of behavior. I illustrate this on two chosen examples where similarity judgments come into play. The first concerns the quantitative similarity in the context of biometric technology. The other, being a case of qualitative similarity, is exemplified by the kinds of judgments that enter the decision-making of which model organism to use for a particular research task.

Similarity is thus an important notion. Notwithstanding the various objections raised against construing similarity as the account of scientific representation, we should not lose sight of that fact.

Bolinska, A. (2013): Epistemic representation, informativeness and the aim of faithful representation. *Synthese* 190 (2), 219-234.

Callender, C. – Cohen, J. (2006): There is no special problem about scientific representation. *Theoria* 55 (1), 67-85.

Chakravartty, A. (2010): Informational versus functional theories of scientific representation. *Synthese* 172 (2), 197-213.

Frigg, R. – Nguyen, J. (2017): Models and representation. In: Magnani, L. – Bertolotti, T. (eds.): *Springer handbook of model-based science*. Dordrecht: Springer, 49-102.

Giere, R. N. (1988): *Explaining science: A cognitive approach*. Chicago and London: The University of Chicago Press.

Godfrey-Smith, P. (2009): Models and fictions in science. *Philosophical Studies* 143 (1), 101-116.

Hughes, R. I. G. (1997): Models and representation. *Philosophy of Science* 64 (Supplement. Proceedings of the 1996 Biennial Meetings of the Philosophy of Science Association. Part II: Symposia Papers), S325-S336.

Knuuttila, T. (2017): Imagination extended and embedded: Artifactual versus fictional accounts of models, *Synthese*, doi: 10.1007/s11229-017-1545-2

Mäki, U. (2011): Models and the locus of their truth. *Synthese* 180 (1), 47-63.

Odenbaugh, J. (2015): Semblance or similarity? Reflections on Simulation and similarity. *Biology & Philosophy* 30 (2), 277-291.

Suárez, M. (2003): Scientific representation: Against similarity and isomorphism. *International Studies in the Philosophy of Science* 17 (3), 225-244.

Thomson-Jones, M. (2010): Missing systems and the face value practice. *Synthese* 172 (2), 283-299.

Toon, A. (2012): Similarity and scientific representation. *International Studies in the Philosophy of Science* 26 (3), 241-257.

Weisberg, M. (2013): *Simulation and similarity: Using models to understand the world*. Oxford and New York: Oxford University Press.