British Society for the Philosophy of Science

Annual Conference
5-7 July 2023
University of Bristol

Wednesday 5th July

10:30-11:30	**BSPS Annual General Meeting (All BSPS members welcome)** (Physics G42)
11.30-12:30	**BSPS Executive Committee Meeting (BSPS Committee only)** (Physics G42)
12:30-14:00	**Graduate Student Workshop with Jacqueline Wallis, Helene Scott-Fordsmand & Adrian Currie** (Physics G42)
14:00-15:00	**Meet the BJPS Editors with Wendy Parker & Beth Hannon** (Physics G42)
15:00-15:30	Coffee (Fry Atrium)
15:30-16:00	**Welcome** (Physics G42)
16.00-17:30	**Plenary Lecture: Peter Vickers** (Physics G42) **"IASC: The Institute for Ascertaining Scientific Consensus"** Chair: Robert Northcott
17:30-19:00	Reception, sponsored by University of Chicago Press (Fry Atrium)

Thursday 6th July

	Stream 1 (Fry G13)	Stream 2 (Fry G09)	Stream 3 (Physics G42)	Stream 4 (Fry LG12)	Stream 5 (Fry G16)
09:30-11:00	Symposium **Function & Dysfunction in Mental Illness** Chair: Rachel Cooper Helen Taylor The Role of 'Dyslexia' in Human Adaptation Justin Garson Function and Dysfunction Paradigms and the Authority of Psychiatry Vaughan Bell & Sam Wilkinson The Mysterious Absence of 'Organic Overlay': Implications for the Conceptual Basis of Disorder and Disability in Psychiatry and Neurology Harriet Fagerberg Psychiatric Disorders as Dysfunctions of Highly Neuroplastic Traits	Symposium **Measurement between the Natural & the Human Sciences** Chair: Corey Dethier Ahmad Elabbar Climate Justice and Measure of Excess Emissions: A Critical Perspective Miguel Ohnesorge & Cristian Larroulet Philippi Is Physical Measurement Relevantly Similar to Human Science Measurement Morgan Thompson Norms for Construct Development in the Social Sciences	Symposium **Machine Learning in Contemporary & Future Science** Chair: William Peden Emily Sullivan Idealization in ML and xAI Andre Curtis-Trudel, Tjonnie Li, William Peden & Daniel Rowbottom Machine Learning and the Problem of Noise-Dominated Measurement Florian Boge Deep-Learning Robustness for Scientific Discovery: The Case of Anomaly Detection Mario Krenn Towards an Artificial Muse for new Ideas in Physics	Symposium **Planet-sized Integrated Philosophy of Science** Chair: Haixin Dang Ann C. Thresher How to Build a Telescope and Community Trust Niels C. M. Martens How to Organise a Telescope Collaboration Juliusz Doboszewski How to Program a Telescope Jamee Elder How to Detect a Photon Ring: Interpreting Images from an Earth-Sized Telescope	Symposium **Perspicuous Representation in Physics** Chair: Samuel Fletcher Henrique Gomes Perspicuous Representation and Perspicuous Understanding of Symmetry-Invariant Structure Caspar Jacobs On the Role of Dynamics in Perspicuous Representation Neil Dewar Against 'Perspicuity' Jill North On the Idea of Perspicuous Representations
11:00-11:30	Coffee (Fry Atrium)				

11:30-13:00

Plenary Lecture: Phyllis Illari (Physics G42)
"What Can Causal Pluralism Do?"

Chair: James Ladyman

13:00-14:00 Lunch (Fry Atrium)

14:00-15:30

Stream 1 (Fry G13)	Stream 2 (Fry G09)	Stream 3 (Physics G42)	Stream 4 (Fry LG12)	Stream 5 (Fry G16)
Symposium **Philosophically Analysing Expert-by-Experience Involvement in Psychiatry** Chair: Riana Betzler Sam Fellowes How the Lived Experience of Experts-by-Experience Relates to the Abstract Nature of Science Lisa Bortolotti, Michael Larkin & Michele Lim Expertise or Perspective in Dialogue? The Role of Lived Experience in the Mental Health Context Astrid Fly Oredsson Academic Affect Norms as a Barrier to Inclusion of Experts-by-Experience in Psychiatric Research	Symposium **Large Scale Brain Models, From Technology to Biology** Chair: Nedah Nemati Tara Mahfoud The Biological Imitation Game: Sublime Explorations of the Boundaries Between Human, Animal and Machine in Large-Scale Brain Modelling Rosa Cao Measuring Similarity in Artificial Neural Networks and Biological Brains Mazviita Chirimuuta 'What I Have not Made, I Do not Understand': Explaining via Remaking of the Brain	Contributed **Physics (Space)** Chair: James Ladyman Ruward Mulder Is spacetime curved? Underdetermination of relativistic gravity theories Joshua Eisenthal Back to the Problem of Space Samuel Fletcher The Definition of Spacetime Singularities, Revisited	Contributed **Explanation & Understanding** Chair: Neil Dewar Franziska Reinhard Elucidating and Embedding: Two Functions of How-Possibly Explanations Tuomas Tahko The Tracking View of Mathematical Explanation Oscar Westerblad What is the nature and function of pragmatic understanding?	Contributed **Values in Science** Chair: Ann C. Thresher Doohyun Sung In Pursuit of Perils: A Social-Epistemological Case Study of Research Method Development in the Biomedical Sciences Jacopo Ambrosj*, Hugh Desmond & Kris Dierickx How are researchers told to deal with non-epistemic factors in science? A content analysis of European national documents on research integrity

15:30-16:00 Coffee (Fry Atrium)

16:00-17:30

Stream 1 (Fry G13)	Stream 2 (Fry G09)	Stream 3 (Physics G42)	Stream 4 (Fry LG12)	Stream 5 (Fry G16)
Contributed: **Biology & Cultural Evolution** Chair: Tim Lewens Pietro Allegretti Sense of Beauty and Aesthetic Predisposition in Evolutionary Aesthetic Theorising	Contributed: **Neuroscience** Chair: Nina Poth Marco Facchin Neural Representations Unobserved Bojana Grujicic Representational similarity analysis underdetermines similarity of object recognition mechanisms in deep neural networks and the brain Nedah Nemati On the Benefits of 'Hand' Engineering in Neuroscience	Contributed: **Physics** Chair: Alastair Wilson Lorenzo Lorenzetti Two Forms of Functional Reductionism in Physics William Wolf* & Karim Thébault Explanatory Depth in Primordial Cosmology: A Comparative Study of Inflationary and Bouncing Paradigms Sarwar Ahmed An Inferential-Information Transmission Account of Observation	Contributed: **History of Philosophy of Science** Chair: Joshua Eisenthal Noah Friedman-Biglin An Apology for Analytic Philosophy: The Left Vienna Circle in Postwar American Academia Marij van Strien On the Value and Challenges of Pluralism in Science: Feyerabend and Bohm in Bristol	Contributed: **Values in Science** Chair: Helene Scott-Fordsmand Tarun Menon & Jacob Stegenga* Sisyphean Science: Why Value Freedom Is Worth Pursuing Stephan Guttinger Transparency and Trust in Science Milena Ivanova Beauty in Experiment: Beyond the Case Study

18.00-20:30 Conference Dinner (The Square Club)

Friday 7th July

	Stream 1 (Fry G13)	Stream 2 (Fry G09)	Stream 3 (Physics G42)	Stream 4 (Fry LG12)	Stream 5 (Fry G16)
09:30-11:00	Contributed: **Biology** Chair: Margarida Hermida Philipp Spillmann How to Explain the Molecular Make-Up of Life on Earth? Tim Lewens The Attractions of Cultural Selection Yihan Jiang The Metaphysics of Mechanisms: An Ontic Structural Realist Perspective	Contributed: **Cognitive Science** Chair: Max Jones Giorgio Sbardolini Skepticism about Common Knowledge Johan Heemskerk Can we extract a theory of content from cognitive science? Oliver Holdsworth An Eliminativist Account of Validity in Psychology	Contributed: **Physics (Quantum)** Chair: Charles Sebens Sebastien Rivat Wait, Why Gauge? Alexander Franklin Everettian Probability as Deterministic Chance Michael Miller Why Go Effective?	Contributed: **Social Science** Chair: Robert Northcott Alexander Linsbichler Reasoning with Models in Thought Experiments: Applying Häggqvist's Template to the Social Sciences Karl Landstrom The epistemic and moral ramifications of Epistemic Extractivism for data-sharing in social science research collaborations.	Contributed: **Values & General Philosophy of Science** Chair: Jacob Stegenga Thijs Ringelberg What is Credit in Science? A Value-Based Interpretation of the Credit Maximisation Approach to the Social Philosophy of Science Corey Dethier When is a Graph Honest? Ethics and Simplification in Science Communication
11:00-11:30	Coffee (Fry Atrium)				
11:30-13:00	**Plenary Discussion: Alexandra Freeman and Marcus Munafo** (Physics G42) **"Creating, Communicating and Evaluating Science in the 21st Century"** Chair: Sabina Leonelli				
13:00-14:00	Lunch (Fry Atrium)				

	Stream 1 (Fry G13)	Stream 2 (Fry G09)	Stream 3 (Physics G42)	Stream 4 (Fry LG12)	Stream 5 (Fry G16)
14:00-15:30	Symposium **The Many Faces of Empiricism: Physical Theories Between Idealism & Realism** Chair: Sebastien Rivat Noah Sterneroff Cassirer and Weyl on Helmholtz (and the Empirical Foundations of Scientific Thought) Philipp Berghofer Realism in Quantum Mechanics? Lessons from Husserl's Empiricism Guy Hetzroni Realism and the Projectibility of Invariance Arguments	Symposium **Major Transitions in Biology, Culture & Cognition** Chair: Pietro Allegretti Samir Okasha The Philosophical Significance of Major Transitions in Evolution Andrew B. Barron, Marta Halina & Colin Klein Major Transitions in the Evolution of Cognition Arsham Nejad Kourki & Ross Pain Major Transitions and Cultural Evolution: Key Challenges and New Directions	Symposium **Similarity & Economic Models** Chair: Öykü Ulusoy Igor Douven Analogical Reasoning: A Carnapian Approach Benjamin Genta Inferring Relations by Analogy Robert Northcott Similarity and Fragility Nadia Ruiz Beyond the Adequacy-for-Purpose View: Model-Building Constraints	Symposium **Quantitative Approaches to Science at Scale** Chair: Kathleen Creel & Liam Kofi Bright Don't Use Machine Learning to Evaluate Grants Haixin Dang & Zina Ward Multi-Analyst Studies and Permissive Evidence	Symposium **Arrows of Time** Chair: Juliusz Doboszewski Bryan Roberts Is Reversibility a Problem for the Reduction of Thermodynamics to Statistical Mechanics? Katie Robertson Discovering Simplicity Dominic Ryder Directed Temporal Asymmetry from Scale-Invariant Dynamics: Lena Zuchowski From Randomness to the Arrow of Time
15:30-16:00	Coffee (Fry Atrium)				

	Stream 1 (Fry G13)	*Stream 2 (Fry G09)*	*Stream 3 (Physics G42)*	*Stream 4 (Fry LG12)*	*Stream 5 (Fry G16)*
16:00-17:30	Contributed **Scientific Theories** Chair: Noah Friedman-Biglin Miklos Rédei & Marton Gömöri Entropic taming of the Look Elsewhere Effect Teodor-Tiberiu Calinoiu Ontic Structural Realism and Accounts of Theories David Wallace Stating Structural Realism: Mathematics-First Approaches to Physics and Metaphysics.	Contributed **Cognitive Science: Causation & Agency** Chair: Max Jones Patrick McGivern Minimal Agency and the Locus of Causation Christopher Joseph An The evolution and ontogeny of normative agency: The role of juvenile social play as behavioural exaptation within a resource-rich developmental niche	Contributed **Physics** Chair: Guy Hetzroni Sean Gryb Counting what counts: symmetry, possibility and inference Charles Sebens Eliminating Electron Self-Repulsion	Contributed **Laws and Analogy** Chair: Milena Ivanova Josh Hunt Expressivism about Laws and Meta-laws Callum Duguid A tension between pragmatic Humeanism and realist metaphysics Helene Scott-Fordsmand* & Mauricio Suárez Negative Analogies and Representation in Medical Practice: A Case from Clinical Orthopaedics	Contributed **Machine Learning & Science** Chair: Nina Poth Non-symbolic few-shot learning Alexander M. Mussgnug Regulative Operationalism

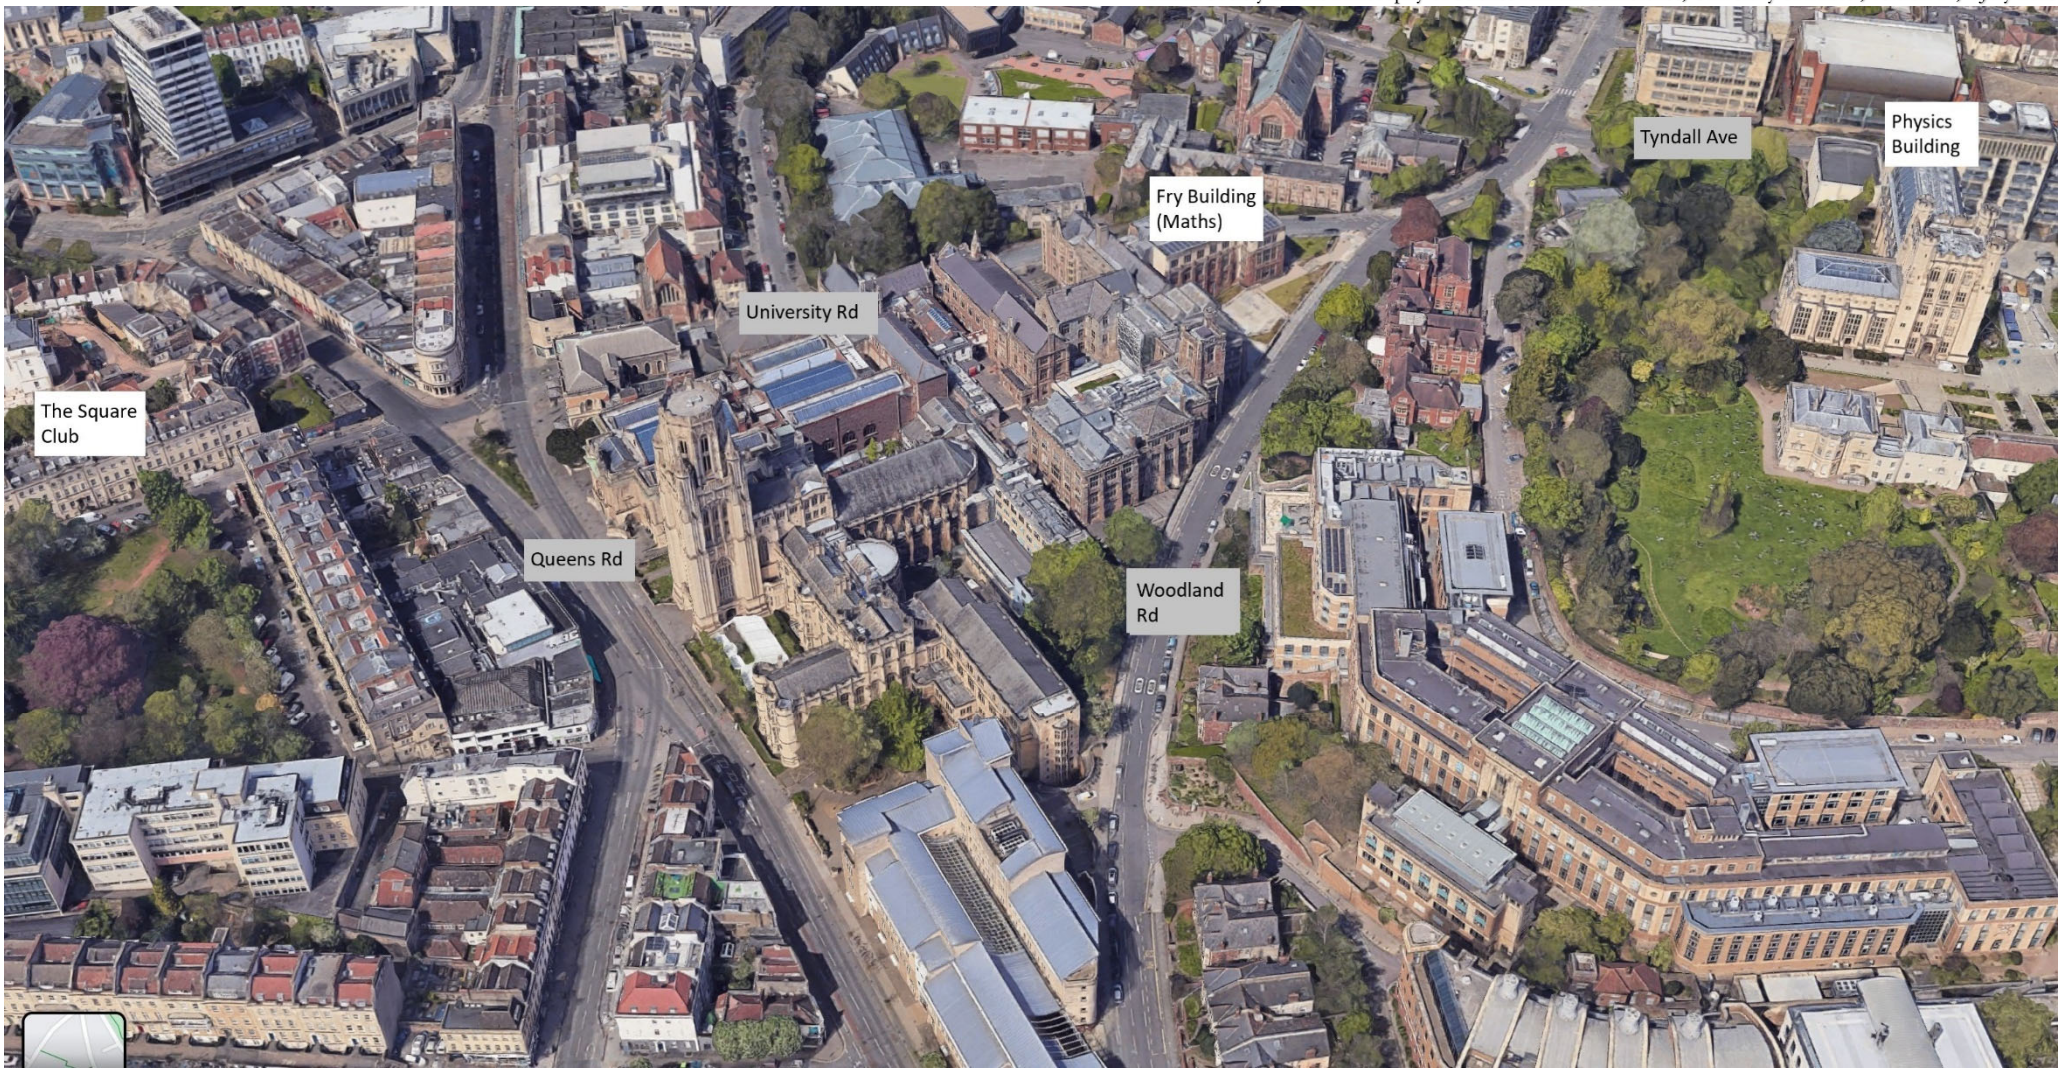

Key Locations:

- **Fry Building (School of Mathematics):** Woodland Road, BS8 1UG
- **Physics Building (School of Physics):** Tyndall Ave, BS8 1TL
- **The Square Club:** 15 Berkeley Square, BS8 1H

General Information

Presenting at the Conference

All rooms have facilities to show slides or other digital presentations. Please talk to the chair shortly before the session to upload your presentation onto the room's computer. This will allow presenters to change quickly and avoid unnecessary delays**.**

Please note that remote presentations are not possible.

Facilities

The majority of the conference takes place on the ground floor of the Fry Building (see map).

Plenary lectures and events for all BSPS members take place in the Powell Lecture Theatre, G.42, in the Physics Building. The Physics Building is about a 400m walk away from the Fry Building and is located on Tyndall Avenue.

There will be signs to the different venues located around campus and the student helpers will be happy to direct you.

Registration

Please check in when you arrive at the conference. The registration desk will be located on the ground floor of the Fry Building.

Conference Dinner

The conference dinner takes place at the Square Club (see map). Participants who registered for the dinner can arrive from 6pm onwards and a buffet of BBQ food will be served from 6:30-8pm. There will be vegetarian, vegan, gluten-free and a limited number of kosher and halal options. Drinks – other than a welcome offering – are not included in the dinner fee and should be ordered at the bar.

Tea and Coffee Breaks

Tea, coffee and finger food will be served in the Fry Atrium at 11:00 (11:30 on Wednesday), 13:00 and 15:30 (15:00 on Wednesday).

The University of Bristol's main food court is located diagonally opposite the Fry Building in Senate House and has a variety of hot and cold food vendors.

Local Organisers

If you have any questions during the conference, please ask Lena Zuchowski or Max Jones (both will be wearing name badges), or one of the student helpers.

Acknowledgments

The BSPS Executive Committee would like to give a warm thanks to the local organisers for hosting this year's conference and to the University of Chicago Press for sponsoring the drinks reception.

From the University of Chicago Press

Journals

The British Journal for the Philosophy of Science

We are honored to publish *The British Journal for the Philosophy of Science* (*BJPS*) on behalf of the British Society for the Philosophy of Science. An international leader in the philosophy of science, *BJPS* showcases outstanding research on a variety of topics, from the nature of models and simulations to mathematical explanation and foundational issues in the physical, life, and social sciences.

Membership in the British Society for the Philosophy of Science includes electronic access to *The British Journal for the Philosophy of Science*.

journals.uchicago.edu/bjps

Also of interest:

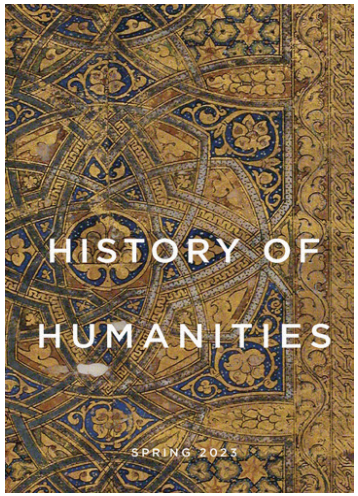

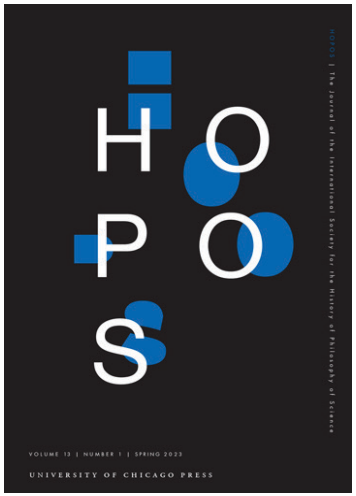

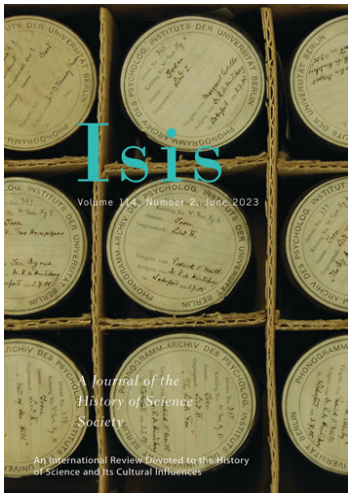

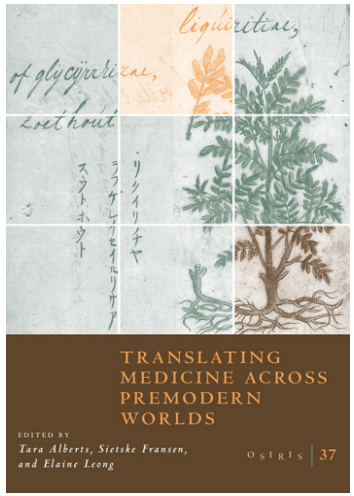

History of Humanities

HOPOS: The Journal of the International Society for the History of Philosophy of Science

Isis: A Journal of the History of Science Society

KNOW: A Journal on the Formation of Knowledge

Osiris 38
Beyond Craft and Code: Human and Algorithmic Cultures, Past and Present

THE UNIVERSITY OF CHICAGO PRESS JOURNALS

journals.uchicago.edu

BSPS Open is an international book series for monographs in the philosophy of science produced in collaboration by the British Society for the Philosophy of Science and the University of Calgary Press.

We believe that scholarly ideas should be freely accessible to everyone, and that scholars should not have to pay to spread these ideas. That is why BSPS Open has adopted a ground-breaking "Open Access Diamond" model that sets a new standard for equality and access. All BSPS Open books are published under a Creative Commons license that allows free publication of the PDF anywhere online, in addition to for-purchase print and digital formats. And, this service is absolutely free to our authors: selections are made only on the basis of scholarly merit, and not the ability to pay Open Access fees. For more information, email bspsopen@thebsps.org or or visit https://www.thebsps.org/bsps-open/

Plenary Lectures

Peter Vickers (Durham University)
IASC—The Institute for Ascertaining Scientific Consensus

A solid, international scientific consensus is the best signal humanity ever gets that a specific scientific statement can sensibly be called an "established scientific fact". Thus, when it comes to the relationship between science and truth, we may think of the scientific realism debate as a distraction, given its preoccupation with first-order evidence (explanatory power, predictive success), as opposed to second-order evidence (consensus). Given the importance of consensus when it comes to reliable, trustworthy, and even 'factual' scientific information, an institute dedicated to measuring scientific opinion efficiently, internationally, and at scale, is well overdue. The 'Institute for Ascertaining Scientific Consensus' pilot project is currently underway, and conducted its first survey in June 2023. It targeted 20,000 scientists across 30 institutions in 13 countries, achieving results in two weeks.

Phyllis Illari (University College London)
What Can Causal Pluralism Do?

I would like to acknowledge a great deal of shared work with my colleague and friend, my philosophy-sister Federica Russo, and my collaborators on the Evaluating Evidence in Medicine project.

In becoming more practice-engaged, philosophy of science has often become more pluralist, and philosophy of causality is no exception to this. One can be at least metaphysically, conceptually, epistemologically, or methodologically pluralist about causality. This leaves the landscape rich and interesting, but also very muddied. I will examine what a causal pluralism that embraces the richly muddy landscape can do. This will involve beginning with a broad conceptual pluralism about causal concepts as all being useful for some jobs, but no account being able to solve all the problems. Second, I argue that we need to add to the library of useful concepts an account of causal production as information transmission, because it succeeds in addressing problems not yet well dealt with, including causal inference methods like marking and tracing. Finally I connect this with the idea of evidential pluralism in medicine, where causal claims are typically justified by evidence of both difference making and of mechanism. Along the way, I hope to convince you that a complicated pluralist picture of causality can be an exciting and fruitful philosophical enterprise.

Plenary Discussion:

Alexandra Freeman (University of Cambridge) & Marcus Munafo (University of Bristol) with Sabina Leonelli (University of Exeter)
Creating, Communicating and Evaluating Science in the 21st Century

A discussion between a scientist and an expert on science communication on how we create and disseminate knowledge in a rapidly changing world. Chaired by Sabina Leonelli, the conversation will cover the scientific process itself (including the role of transparency), and novel approaches to scholarly communication.

*

Symposium Abstracts

Function and Dysfunction in Mental Illness

We propose a symposium on the topic of function and dysfunction in mental illness. The history of psychiatry evidences a clash between two distinct models for understanding psychopathological phenomena. One the first model, mental illness indicates dysfunction: there has been a breakdown of a mechanism or part within the person, which now requires 'fixing'. On the other model, so- called mental illnesses are often functional or goal-directed responses to life stressors or adverse environmental conditions. Our symposium will bring together leading scientists and philosophers to present interdisciplinary research on the role of functional and dysfunctional processes in mental illness. We will consider the unique properties of the brain, and the relationship between mental functions and neural processes, and the status of specific mental disorders such as dyslexia. We will also consider the implications of a function-based model for the legitimacy of psychiatry as a medical discipline.

Helen Taylor (University of Cambridge, University of Strathclyde)
The Role of 'Dyslexia' in Human Adaptation

Research on dyslexia has primarily focused on educational difficulties, with theories framing differences in neurocognitive processes as deficits. Consequently, people with dyslexia are variously described as having a learning difficulty or neurobiological disorder. This long-standing deficit-centric view provides an incomplete picture, however. Even the very earliest accounts of dyslexia noted the frequent occurrence of countervailing strengths in areas such as spatial and non- verbal reasoning and evidence of a 'dyslexic advantage' has since been accumulating. The two purposes of this presentation are to a) provide an alternative view of dyslexia-associated cognition that explains the pattern of both strengths and difficulties from an evolutionary perspective, and b) explain this new understanding in a broader context, in order to realize the significance that this reframing may have for understanding the process of human adaptation and cultural evolution. The research that will be presented proposes that the various cognitive strengths discovered or proposed in people with dyslexia can all be seen as various facets of a cognitive specialization in more global exploratory learning, which appears to confer particular abilities in understanding of complex system dynamics. The presentation will draw on various disciplines to consider both theoretical and qualitative evidence for the role this plays in human adaptation. In particular it will discuss how global exploratory learning acts in complementary fashion to other learning strategies that excel in more local level exploration as well as exploitation of existing knowledge. Finally the implications will be discussed in an evolutionary framework and in the context of modern society, in terms of our approach to educational practices, organizations and sustainability, including the importance of balancing this system.

Justin Garson (City University of New York)
Function and Dysfunction Paradigms and the Authority of Psychiatry

Since the 1970s, American psychiatry, and to some extent global psychiatry, has been largely in the grips of a paradigm I call "madness-as-dysfunction." In this view, when someone is mad, it is because something has broken inside of them; the purpose of the psychiatrist is to fix it. In my historical work, I've recovered an alternative paradigm which I call "madness-as-strategy." In this paradigm, madness is seen as having a purpose, goal, or function. I'll illustrate this latter view with recent examples drawn from the study of depression, personality disorders, and psychosis. These examples suggest that we may be currently witnessing a paradigm change from madness-as-dysfunction to madness-as-strategy. The purpose of this presentation is to pose a simple question: to what extent is madness-as- dysfunction essential, rather than incidental, to psychiatry? It's natural to think that psychiatry could survive a "paradigm shift" from madness-as-dysfunction to madness-as-strategy; historians tell us that it undergoes such paradigm shifts all of the time. On the other hand, psychiatry is, by definition, a branch of medicine; this is what distinguishes psychiatry from other mental health professions. Hence, at least implicitly, psychiatry, as such, is committed to the claim that the kinds of distressing or disturbing mental phenomena that it treats are, in fact, medical conditions, that is, that they fall under the jurisdiction of medical professionals. But what must mental disorders be in order to fall under the jurisdiction of medical professionals? I argue for the following answer: to warrant the attention of medical professionals, mental disorders must be pathologies – that is, that they must be caused, in part, by inner dysfunctions. If mental disorders must be pathologies in order to warrant the attention of medical professionals, then that suggests that psychiatry, as such, is committed to madness-as-dysfunction. To challenge madness-as-dysfunction, then, is not to call for a "paradigm shift" within psychiatry. Rather, it is to challenge psychiatry as such.

Vaughan Bell (University College London, South London and Maudsley NHS Trust) & Sam Wilkinson (Exeter University)
The Mysterious Absence of 'Organic Overlay': Implications for the Conceptual Basis of Disorder and Disability in Psychiatry and Neurology

"Functional overlay" is a concept in psychiatry and neurology where impairment or disability resulting from confirmed pathophysiology, typically neurological disorder, is judged to be greater than the neurological disorder itself can account for – indicating that 'functional' factors apparently 'overlay' the neurological disorder to account for rest of the disability. But there is no concept of "organic overlay", where the extent of disability must be accounted for by the presence of an additional layer of pathophysiology that 'overlays' a 'functional' disorder or factors. There are several

examples that would seem to fit a concept of "organic overlay" – where 'organic' problems arise from 'functional' factors and increase impairment or disability. We discuss several well-confirmed examples in neuropsychiatry: reflex epilepsies – where seizures are triggered by experiences, thoughts or emotions; bladder damage after functional urinary retention; and cataplexy – a common symptom of narcolepsy where sudden muscular weakness is triggered by strong emotions. Given this, we argue that the status of "functional overlay" and the absence of "organic overlay" reveals tacit conceptual components of causation in medicine. Namely, that there is a presumed hierarchy of causation from 'organic' to 'functional'. The concern with marking out patients who show 'functional additivity' but not 'organic' additivity indicates that organic factors act to mark the outer boundary of unquestionably legitimate disorder in medicine, likely reflecting lingering moral judgements about causes of disability. And that the conceptual basis of clinical practice is incompatible with dynamic, and likely more scientifically accurate, models of neuropsychiatric disorder that theorise that disorders and disabilities arise from an interaction between different levels of explanation.

Harriet Fagerberg (City University of New York, King's College London)
Psychiatric Disorders as Dysfunctions of Highly Neuroplastic Traits

This paper proceeds from the assumption that mental disorders are brain disorders. Contra the received view, mental disorders are not distinct from brain disorders like software bugs are distinct from hardware malfunctions in a classical computer. However, even granting that mental disorders are brain disorders, there may yet be systematic differences between psychiatric disorders and somatic disorders owing to the *kinds* of neural traits which are affected in each case. This paper explores the hypothesis that psychiatric disorders are special because they are dysfunctions of highly neuroplastic traits. I start is with the observation that psychiatric disorders often involve disturbances in higher- level or complex neural functions (fear attribution, habit formation, mood regulation and so on) which are responsive to selection, compensation and adaptation. Psychiatric disorders are, as the National Institute of Mental Health's Research Domain Criteria (RDoC) Project would have it, "neural circuit disorders". Paradigm somatic disorders, in contrast, are characterised by disturbances in relatively rigid mechanisms wherein in little plastic change is possible, or where the degree of responsiveness is 'hard-wired'. I give an account of what it takes for a neuroplastic trait to be dysfunctional, and consider how dysfunctions of plastic mechanisms might 'behave' differently to dysfunctions of rigid mechanisms. I build on this in clarifying RDoC's conception of psychiatric disorders as disorders of neural circuits, laying the foundations for an account of the evolved functions of higher-level neural circuits and systems. On this basis, I suggest that psychiatric disorders can be fruitfully understood as higher-level brain dysfunctions wherein no *part* is dysfunctional. This model can accommodate the notion that psychiatric disorders are genuine biological dysfunctions, and as such 'real' disorders, while simultaneously accounting for the relative lack of success of the reductionist programme in biological psychiatry.

Measurement between the Natural and the Human Sciences

Quantitative measurement drives research in the natural sciences. Yet philosophers and scientists remain divided about the extent to which quantitative measurement can and should be applied in the human sciences. Particular worries concern the causal complexity of social and psychological phenomena and the moral and political consequences that their measurement can have. In this symposium, we develop a new perspective on these problems by considering the measurement of *concepts that lie at the intersection of the natural and human sciences*. In doing so, the papers in our symposium collectively explore the extent to which insights from the philosophy of measurement in the natural sciences can be applied to the human sciences. Our outlook is both practice-oriented and systematic. To that end, we have chosen three contributions that develop detailed case studies: the first examines the measurement of "biodiversity" in ecology, the second explores the quantification of "fertility" in biomedicine, and the third focuses on "excess emission" measurement in climate economics. All three measurement practices are widely recognised to be deeply value-laden and to involve high definitional and procedural uncertainty. Our fourth contribution takes a step back and asks systematically which (if any) analogies between physical and human science measurements are methodologically salient, i.e., what (dis-) similarities between them offer relevant insights for the prospects of measurement beyond the natural sciences.

Ahmad Elabbar (University of California San Diego)
Climate Justice and Measures of Excess Emissions: A Critical Perspective

Carbon accounting, the practice of measuring the greenhouse gases emitted by various actors in the global economy, plays a crucial role in the politics of climate change. Both at the state level and internationally, carbon accounts constitute the empirical record against which promises of mitigation are assessed, and the basis for apportioning moral and legal responsibility for climate loss and damage (see Steininger et al. 2016). Traditional carbon accounting methods, such as those recognised and used in UN climate accords, constitute what we might call 'thin' measures of carbon emissions: they aspire to characterise their *measurands* – the emissions of, say, nation states – in purely descriptive terms, eschewing evaluative concepts. Familiar numbers and charts representing the carbon footprints of nations are the result of these thin measures (e.g., IPCC 2014, chap. 5). Recently, however, a group of climate economists and sustainability researchers have objected to the use of traditional methods of carbon accounting on grounds of climate justice (Hickel 2020; Hickel et al. 2022). They argue that traditional methods, despite their diversity, all fail to measure the morally relevant

portion of an agent's emissions. The alternative they propose is a 'thick' measure of 'excess emissions', a measurand characterised explicitly in moral terms: drawing on conceptions of global distributive justice and historical responsibility. Their analysis of excess emissions reveals stark global inequalities in energy use across nations that are obscured by traditional carbon accounting which suggest that the contribution of the Global South to climate change rivals or exceeds that of the Global North. On their measure of excess emissions, by contrast, the Global North is responsible for 92% of morally relevant emissions. My aim in this presentation is to consider what it would take to justify such thick measures of carbon emissions, considering whether the particular proposal of excess emissions developed by Hickel et al. (2020; 2022) meets this challenge. I begin on a familiar note, pointing out that thick measures raise unique problems of justification that do not arise for thin measures, namely the justification of the evaluative choices made in characterising the measurand. Drawing on recent work on measurement in the social sciences (e.g., Schroeder 2019; Murray and Schroeder 2020; Alexandrova and Fabian 2022), I consider two broad approaches that might be pursued: a substantive approach, where the moral assumptions necessary for characterising excess emissions are justified in terms of first-order moral reasons, and a pure proceduralist approach, where the relevant moral assumptions are fixed via a legitimate political procedure. I argue that Hickel et al.'s proposal falls short of providing a justification on either strategy, and that the prospects of providing such justifications are slim, given deep theoretical and practical challenges in the ethics and politics of climate change. Despite this negative assessment, I conclude by considering the positive role figures of excess emissions can play in climate change discourse, provided their limitations qua measurements are acknowledged and they are used to open, rather than close, conversations over responsibility for climate breakdown.

Miguel Ohnesorge (University of Cambridge) & Cristian Larroulet Philippi (University of Cambridge)

Is Physical Measurement Relevantly Similar to Human Science Measurement?

There is a long-standing and unresolved dispute about the possibility of quantitative measurement in the human sciences. Both optimists and pessimist in this debate justify their conclusions by highlighting similarities and dissimilarities to quantitative measurement in physics. We show that the development of the debate has been driven by evolving views about which such similarities must obtain to make quantitative measurement possible. Initially concentrating on dissimilarities to direct physical extensive measurement operations (Campbell 1928), the debate gradually turned to similarities to physical conjoint measurement operations (Luce and Tukey 1964; Krantz 1964; 1991; Schwager 1991; Michell 1997; Sherry 2011). Recent disagreements about the applicability of conjoint measurement, finally, focussed on dissimilarities to the experimental control over measurands and confounders in experimental physics (Trendler 2009; 2019a; 2019b; Michell 2019; Krantz and Wallsten 2019). We think that the evolution of this debate has been broadly informative, in so far that it brought more relevant similarities to the attention of scientists and philosophers. As it stands, dissimilarities in experimental control provide good reasons for being pessimistic about quantitative measurement in the human sciences. Our aim is not just descriptive, however. After recapturing the role of analogy in the debate, we argue that the class of relevant similarities so far considered remains incomplete and that this incompleteness undercuts the pessimistic conclusion. As of now pessimists have only established dissimilarities to allegedly representative measurements in experimental physics, such as Galileo's inclined plane measurements of acceleration in free fall (Trendler 2009). We discuss the measurement of seismic moments to show that the history of physics also contains successful measurements in which scientists could neither experimentally control their stipulated measurand nor any confounding factors. Seismologists compensated for their lack of experimental control by relying on a methodology of theory- and model-mediated measurement, in which they gradually modified modelling assumptions about the measurand, and confounders based on their ability to lead to (i) convergent, (ii) stable, and (iii) increasingly precise measurement outcomes (Howell 2005; Smith 2007; Miyake 2011; 2017b; 2017a). The success of this methodology - culminating in the Global Seismographic Network and the double-couple model of seismic sources - shows that experimental control is not a necessary condition for quantitative measurement. Our conclusion does not by itself justify optimism about human science measurement but refutes the strongest argument against its possibility. Moreover, it points to a class of physical measurements akin to seismology, which contain more relevant similarities to human science than the narrow class of cases considered up to now. While we primarily aim to move forward the existing dialectic about the possibility of human science measurement, seismology also contains another unconsidered similarity to psychological measurement: the role of non-epistemic values. For long, alternative definitions and operationalisations of earthquake "intensity" or "magnitude" coexisted because of their respective ability to fulfil different, morally significant purposes (e.g., protecting buildings vs detecting likelihood of tsunamis vs promising earthquake prediction). Notwithstanding such long-standing value trade-offs, seismologists made scientific and, eventually, practical progress by isolating a single scalar measurand (seismic moment)

Morgan Thompson (Bielefeld University)

Norms for Construct Development in the Social Sciences

Racial discrimination is a multi-dimensional construct (or "Ballung" concept; Cartwright, Bradburn, and Fuller 2016) as are many constructs in social science. It can be interpersonal and/or institutional as well as clearly attributable to one's race or ambiguously attributable. In public health research, all of these dimensions contribute

to explanations of racial health disparities (Thompson 2023a). Yet, when public health researchers develop subjective measures of racial discrimination, the measures all focus on the same dimensions. While measurement often requires narrowing focus to only particular aspects of multi-dimensional constructs (e.g., Chang 2004), it can be detrimental when all measures focus on the same set of dimensions (call it path dependence). In previous work, I argued that this path dependence in measurement cannot be prevented by adopting a coherentists view of measurement (Thompson 2023b). The main take-away is that norms are needed to evaluate the initial refinement of the construct rather than taking it for granted. Here I turn from the explicating the problem to proposing norms for the development of constructs (with reduced dimensions) for the purposes of measurement. These norms must respect that constructs are often based on folk concepts (e.g., Bloch-Mullins 2022) and social theoretical concepts. Questions to be addressed include: when should qualitative dimensions be dropped from the construct? What are the ethical and epistemic limitations in revising dimensions of the folk or social theoretical concept? And how can the risk of path dependence in measure development be mitigated?

Machine Learning in Contemporary and Future Science

AI techniques are increasingly used in science, with striking and remarkable results. Yet philosophers of science are only beginning to grapple with this development as it pertains to contemporary science and its ramifications for future science. This symposium will bring together philosophers of science and working scientists to address some of these issues. It will pay particular attention to how AI techniques are used in various scientific subdisciplines, including radio astronomy and gravitational wave astrophysics, and the resultant challenges for future science in these areas. Topics to be considered include supervised vs unsupervised AI techniques for scientific discovery, how AI is being used to complement and supplant citizen scientists, and parallels between ML and traditional modelling techniques.

Emily Sullivan (TU Eindhoven)
Idealization in ML and xAI

Interpretability and xAI methods are important for establishing trust in using black-box models. However, recently criticism has mounted against current xAI methods that they disagree, are necessarily false, and can be manipulated, which has started to undermine the deployment of black-box models. Rudin (2019) goes so far to say that we should stop using black-box models altogether in high-stakes cases because xAI explanations 'must be wrong'. However, strict fidelity to the truth is historically not a desideratum in science. Idealizations--the intentional distortions introduced to scientific theories and models--are commonplace in the natural sciences and are seen as a successful scientific tool. Thus, it is not falsehood qua falsehood that is the issue. In this talk, I outline the need for ML and xAI research to engage in idealization evaluation. I discuss where current research can help with idealization evaluation and where innovation is necessary. I address questions surrounding how idealization in highly idealized models differ from idealizations deployed in ML and how ML idealizations can aid scientific inquiry.

Andre Curtis-Trudel (Lingnan University), Tjonnie Li (KU Leuven), William Peden (Lingnan University) & Darrell Rowbottom (Lingnan University)
Machine Learning and the Problem of Noise-Dominated Measurement

The detection of gravitational waves is among the most striking scientific successes in recent years (Abbott et al 2015). One of the primary challenges faced by researchers attempting to detect a gravitational waveform is the inherently noisy nature of the available data. Researchers employ a variety of techniques to extract a target signal from the noisy background. These techniques are informed by background physical theory, which provides crucial information about the expected shape of a target signal. Yet this way of dealing with noisy data raises challenges for the prospect of novel discoveries and breakthrough science. In particular, it becomes difficult to see how one might identify novel phenomena in such data given the role of current theory in filtering out the noise. We call this challenge "the problem of noise-dominated measurement". This paper investigates one aspect of the problem of noise-dominated measurement in more detail: the use of machine learning (ML) techniques to detect a target signal. First, we argue that the use of such techniques exacerbates the problem of noise-dominated measurement. ML systems are often trained, via supervised learning regimens, to look for quite specific signals. Accordingly, such systems are not in a position to distinguish novel and potentially interesting signals from background noise. Second, we consider whether unsupervised learning methods might help to address the problem. Our assessment is mixed. On the one hand, unsupervised methods are generally able to identify novel structures and patterns in data. On the other hand, it is generally difficult to interpret the outputs of such methods without an appropriate conceptual framework (Boge 2021, Kieval forthcoming). Yet these frameworks are typically absent in in breakthrough science. Whether machine learning can provide a way forward thus remains to be seen.

Florian Boge (University of Wuppertal)
Deep Learning Robustness for Scientific Discovery: The Case of Anomaly Detection

Machine Learning (ML) techniques such as Deep Neural Networks (DNNs) are of great promise in science today. In High Energy Physics in particular, they are supposed to foster scientific discovery through the detection of anomalies, without reliance on any specific theory or model. Anomalies, in turn, have long been recognized as a major

driving force of science: Kuhn ([1970]) held them responsible for paradigm shifts, Lakatos ([1970], [1976]), in severe cases, for the abandonment of a theory's hard core, Laudan ([1977]) for the establishment of a preference order among rival theories, and even recent proponents such as de Regt ([2020]) for the advancement of science through an increase of understanding. However, DNNs also have astonishing shortcomings, as they are vulnerable to 'adversarial examples'; data instances that are easily classifiable for humans but totally misclassified by DNNs. Adversarial vulnerability is a double-edged sword: On the one hand, it shows that discerning DNNs' credible outputs from flukes requires some skill. On the other hand, adversarials exhibit DNNs' sensitivity to subtle, often humanly-inscrutable features that could also be scientifically productive (Buckner [2020]). Such features are, in fact, being utilised in anomaly detection. Against this backdrop, I offer an analysis of, and a cautionary tale about, DNNs' present utility for scientific discovery in the talk. To do so, I will introduce a notion of performance robustness, which DNNs need to satisfy in order to be able to deliver genuine discoveries. Furthermore, I will argue that the achievement of performance robustness often, if not always, implies limitations to fully ML-driven discovery.

Mario Krenn (Max Planck Institute for the Science of Light)
Towards an Artificial Muse for new Ideas In Physics

Artificial intelligence (AI) is a potentially disruptive tool for physics and science in general. One crucial question is how this technology can contribute at a conceptual level to help acquire new scientific understanding or inspire new surprising ideas. I will talk about how AI can be used as an artificial muse in quantum physics, which suggests surprising and unconventional ideas and techniques that the human scientist can interpret, understand and generalize to its fullest potential.

Planet-Sized Integrated Philosophy of Science

In 2019, the Event Horizon Telescope Collaboration released the first ever picture of a black hole, created from data collected by using a virtual, Earth-sized telescope comprising dishes on six sites dispersed across the world. The Next Generation Event Horizon Telescope Collaboration aims to increase the number of sites to ~ 20. The History, Philosophy, & Culture Working Group (HPC) constitutes one of its 11 working groups, consisting primarily of scholars from the humanities and social sciences, thereby providing an unprecedented opportunity for these fields to contribute in real-time to a large (astro)physics collaboration. Conversely, it provides a contemporary case study for various approaches, methods and tools within these fields. The HPC has prioritised four topical areas of contribution: 1) *Responsible siting:* taking into account social, cultural, ethical and environmental specifications when choosing telescope sites; 2) *Collaborations*: implementing governance structures that stimulate the desired social epistemology, such as utilising the value of dissenting opinions; 3) *Algorithms, Inference, and Visualisation*: using tools from philosophy and art history to ensure the long term reliability of inferences made via algorithms and simulations, as well as intentional image presentation; 4) *Foundations*: using tools from the philosophy of science to bridge the gap between theory and observation.

Ann C. Thresher (Stanford University)
How to Build a Telescope and Community Trust

Telescope siting has, historically, relied entirely upon ensuring sites meet the technical specifications required for observation including weather, atmospheric clarity, accessibility, and cost of developing the site. Astrophysicists have, however, obligations beyond merely the technical. As the field is becoming increasingly aware, telescopes exist within a broader social context, and the choice of site needs to be responsive to ethical, cultural and environmental specifications just as much as technical ones. We need look no further than the ongoing conflict surrounding the Thirty-Meter Telescope on Mauna Kea or the Square Kilometer Array in South Africa to see the importance of this lesson. As a step forward, the ngEHT is working to develop new guidelines for the responsible siting of telescope projects. As part of this effort a dedicated focus group has been formed and integrated into the site selection process with the goal of examining the impacts of construction and operation at possible sites. These aspects will be a critical part of the decision process, and will include consideration of the environmental impacts alongside broader community and cultural ones. In this, the ngEHT aims to model better ethical research practices in telescope siting and ongoing community collaboration, as well as develop explicit guidelines that can help with future siting challenges. Here, we present the ethical challenges of telescope siting, including considerations of environmental and community impact. Drawing on work in parallel fields including genetic modification, paleontology, and conservation we advance a set of ethical best-practices for astronomy, and consider the fundamental philosophical challenges of building large-scale scientific experiments, particularly in areas of the world with a history of colonial, environmental, and scientific exploitation.

Niels C. M. Martens (Utrecht University, University of Bonn)
How to Organise a Telescope Collaboration

The process and results of knowledge formation within a scientific collaboration, i.e. its social epistemology, depend on its governance and social structure. Large scale scientific collaboration can take place within a variety of governance/organisational structures, ranging from top-down hierarchical structures to more loosely organised bottom-up collaboration in the absence of a formal governing structure. Extreme top-down collaborations prioritise the aims of the collaboration as a whole over the interests of its individual members, implicitly and incorrectly assuming that these are in tension

with one another. On the other hand, too loose a bottom-up approach will be insufficient as sustained collaboration is required for the ngEHT to build the additional telescopes and coordinate the whole process of observing and data analysis to arrive at high- quality black hole videos. Moreover, utilising the epistemic value of diverse opinions and optimising individual creativity is not a mere matter of the absence of top-down governance structure, but also the implementation of positive measures that bring out these epistemic advantages. The main goal of the collaborations focus group is to provide the ngEHT collaboration with advice on optimising its social epistemology by developing a governance structure that is located somewhere in the middle of this spectrum, in a way that is the best of both worlds. Our main starting point is the realisation that supporting the individual—via a collaboration structure that enables, encourages, supports and emphasizes transparent decision-making, diversity, fair credit assignment and accountability—is to the ultimate benefit of the collaboration as a whole, as well as its epistemic outputs. This talk explores the following four topics within the context of the ngEHT collaboration, from the perspectives of social epistemology and the philosophy of scientific practice: 1) governance and team structures, 2) consensus vs. dissenting opinions, 3) a forecasting tournament, and 4) authorship practices.

Juliusz Doboszewski (Harvard University)

How to Program a Telescope

Experiments such as the EHT and ngEHT heavily rely and will continue to rely on computational methods at various stages of its operation. This, in particular, includes image reconstruction algorithms and parameter extraction. The EHT and ngEHT probe previously inaccessible physical regimes, and so for many purposes they provide the only available line of evidence. What makes these methods trustworthy and free from experimenter's (or, rather, programmer's) bias? A partial answer is provided by a robustness-type argument: convergence among independent methods increases the security of the evidence claim. We will compare this form of robustness with robustness understood as variation among models. A related question concerns the sense in which the EHT imaging algorithms might be seen as theory neutral. It turns out that some aspects of the experiment are moderately theory laden: parameter extraction relies on general relativistic magnetohydrodynamic simulations. (Some of the resulting reliance, however, might be limited using multi-messenger techniques.) The above answers concerning trust in the (ng)EHT computational methods, however, lead to caveats. Philosophy of science work on computer simulations and artificial intelligence is helpful in diagnosing some of the long term challenges of fields relying on such methods. The relevant caveats include (i) generative entrenchment – the model might become path dependent, in that order of implementation of various physical effects and overall history of its development leaves an imprint on the model; moreover, awareness of modeler's ad hoc choices might become lost over time; (ii) confirmation holism – assignment of success or failure to a particular component of the module might be hard or even impossible; and (iii) opacity – inner workings of a model may become inaccessible to its users. We will explore both positive and negative aspects of these features, show (using the 2022 Sagittarius A* image) the sense in which they already arise in the EHT, and follow that by sketching some procedures for diagnosing their occurrence and limiting their negative consequences. This is an interesting and novel form of philosophy of science highly integrated with scientific practice.

Jamee Elder (Harvard University, University of Bonn)

How to Detect a Photon Ring: Interpreting Images from an Earth-Sized Telescope

General relativity predicts that images of optically thin accretion around a black hole will contain a "photon ring", a nested series of increasingly sharp subrings from increasingly strongly lensed emission in the region. These are indexed by the number n of half orbits around the black hole, so the n = 0 image is the primary "direct" image (which may not form a ring) and n = 1 is the secondary image formed by photons that have completed a half orbit before reaching the observer. From there, each subsequent ring is both narrower and dimmer. The features of the subrings (n = 1 onwards) are determined by the spacetime curvature. This means that measurements of the photon ring with the ngEHT could provide a clean probe of the target black hole's properties and tests of the Kerr metric [Tiede et al., 2022]. However, there has been recent debate among astrophysicists concerning the prospects for detecting the photon ring; [Broderick et al., 2022] claim to have measured the n = 1 subring for M87* using a method of "hybrid imaging" from EHT data, while [Tiede et al., 2022] advocate for "extreme caution" when using such techniques to interpret both EHT and prospective ngEHT data. From a technical standpoint, it is a matter of debate what would be required to measure or detect the n = 1 ring (e.g., number and placement of telescopes, models, methods). From a philosophical perspective, the case of the photon ring raises several questions about detection, measurement, and evidence. In this talk, I will consider what would (or should) count as a successful detection of the n = 1 ring. I will argue that the challenges for claiming a detection in this case are bound up with the theory- or model-ladenness of the techniques needed to claim a detection; in particular, concerns about false positives are a classic way in which theory-ladenness can undermine confidence in a purported detection. I will compare and contrast the challenges of separating out the signal from the "noise" across this and other case studies from physics and astrophysics, including (for example) LIGO-Virgo and ATLAS/CMS detections. Overall, the case of the photon ring exemplifies both the opportunities and challenges that await the ngEHT Collaboration as it aims to perform new tests of fundamental physics using a virtual Earth-sized telescope.

Perspicuous Representation in Physics

The topic of this symposium is the notion of *perspicuous representation*: representations that are, in some sense, more 'natural' or 'direct'. It is often claimed that perspicuous representations are a *desideratum*, even a *sine qua non*, of theory interpretation. It is often unclear, however, what *makes* a representation perspicuous - if there even is an unequivocal criterion. The papers of this symposium offer various (sometimes competing) accounts of perspicuity, highlighting important applications of that notion. The first paper proffers a definition of perspicuous representations as those that are axiomatically built up from invariant concepts, and applies this to the debate around 'internal' vs 'external' sophistication. The second paper focuses on the role of dynamics: through a case study of inertial frames, it argues that perspicuous representations are those that mesh well with a theory's symmetries. The third paper defends the thesis that there is no non-pragmatic notion of perspicuity. Rather, perspicuous representations are those that aid subjective understanding. The final paper argues that certain representations are perspicuous in a non-pragmatic, objective sense, and applies this to the topic of theoretical equivalence. Together, these papers represent a significant step in a clearer characterisation of perspicuity, with many avenues for further applications.

Henrique Gomes (University of Oxford)

Perspicuous Representation and Perspicuous Understanding of Symmetry-Invariant Structure

In the long-standing debate about symmetry and equivalence, the prevailing modern stance is labelled `sophistication'. This stance requires us to accept models of the theory that are not symmetry-invariant, and yet to hold that these models represent the same physical situation of a given target system. Since the models don't wear the symmetry-invariant structure on their sleeves, we are `motivated' to find more *perspicuous* representations of that structure. But what is perspicuity, exactly? Although theoretical representation can have many purposes, questions about whether a given representation is `perspicuous' are most salient in the discussion of symmetries. In this context, Møller-Nielsen (2017)'s characterises a perspicuous representation as one "which corresponds to, or "limns", reality's structure in some suitably faithful way". This is a loose characterisation, and I will try to fill it in. Under my more regimented interpretation, I will argue that general relativity and gauge theory---the cornerstone of modern physics---have no unique perspicuous representation (p-representation) of symmetry-invariant structure. In more detail, p-representations are required to be concrete and local, and that implies we cannot uniquely p-represent the symmetry-invariant structure of the models of general relativity and gauge theory. This is somewhat conciliatory with Dewar's pragmatic view of perspicuous representation. On the other hand, I will define a symmetry-invariant structure to be perspicuously understood if it admits at least one axiomatic formulation using a basis of metaphysically clear postulates, that are invariant/relational/structural; and I will show that both gauge theory and general relativity admit such a formulation.

Caspar Jacobs (University of Oxford)

On the Role of Dynamics in Perspicuous Representation

The laws of Newtonian mechanics in their coordinate-based form (e.g. F = ma) only hold true (insofar as they are true) with respect to a privileged class of coordinates: the inertial ones. It is often suggested that these are just the coordinates that best, or most perspicuously, represent the theory's spacetime structure. Maudlin (2012) puts it as follows: "The intrinsic geometrical structure of space and time according to Newton entails that special sets of coordinates exist. [...] the existence of such convenient coordinates [...] follow[s] from the spacetime structure itself." Similar sentiments are expressed in Earman (1989), Pooley (2013) and North (2021). I dispute this claim. On my view, *any* bijection between the spacetime manifold and R4 represents that spacetime's structure equally well under *some* representational convention. To put the point differently: claims that certain coordinate systems better reflect spacetime structure than others *presuppose* a particular convention about how certain structural features are represented numerically. But the only (non-dynamical) reason to prefer one such convention over another is for pure convenience. This leaves us with a problem: coordinate-based formulations of spacetime theories presume *some* way to define a privileged class of coordinates, but this does not follow from the structure of spacetime itself. Of course, one option is to adopt a form of conventionalism, but I believe that there is a better solution. In particular, I claim that the laws themselves help to constrain the class of perspicuous representations. Recall that the laws of Newtonian mechanics are Galilean-invariant: their form remains the same under certain transformations of the spacetime coordinates, such as translations and boosts. It thus seems reasonable to require that the ways in which a coordinate system encodes spacetime structure *also* remain the same under these transformations. If this weren't the case, then, given a representational convention, Galilean transformations would seem to make a physical difference, contrary to the Galilean-invariance of the dynamics. I show that in the case of Euclidean space, these are just the standard Cartesian coordinates (of which inertial coordinates are a subclass). It is therefore possible to recover the intuitive claim that these coordinates most naturally represent the structure of space. But *contra* Maudlin, this does not follow "from the spacetime structure itself"; rather, it follows from spacetime structure *and* the theory's dynamics. If the theory's laws were not Galilean-invariant, the most perspicuous coordinates would differ, too. In a slogan: perspicuous representation is not kinematical but dynamical.

Neil Dewar (University of Cambridge)

Against 'Perspicuity'

What is it for a representation to be "perspicuous"? One of the first uses of this terminology in the recent literature occurs in Wallace and Timpson (2010), where a perspicuous representation is described as being "a representation which makes manifest the structure that the theory ascribes to the world". Much of the recent debate, however, draws on Møller-Nielsen (2017)'s characterisation of a perspicuous representation as being one "which corresponds to, or "limns", reality's structure in some suitably faithful way". This is an important shift: from perspicuity as something which makes a representation good for us (by making our commitments manifest) to something which makes a representation intrinsically better (by better depicting the structure of reality). This contrast is noted clearly by North (2021), who defends the latter way of thinking about perspicuity; and even asks, "What do they [Wallace and Timpson] mean by a perspicuous representation, if not a particularly clear-eyed representation of the true nature of the physical world?" This paper seeks to answer that question, and to defend the resulting account of perspicuity as a subjective rather than objective virtue. To do so, I argue that perspicuous representations are those which best afford an understanding of the theory in which they occur—where the relevant notion of understanding is de Regt's pragmatic understanding (de Regt 2009, 2017), i.e., the capacity to effectively use the theory. I bring this analysis to bear on some of the paradigmatic case studies of perspicuity: coordinate-based versus coordinate-free presentations of geometry; phase-space versus normal-space representations of multi-particle systems; and Galilean spacetime versus an equivalence class of Newtonian spacetimes. I show that we can use de Regt's analysis to explicate the sense in which, in each of these pairs, the latter representation is more perspicuous than the former. However, this characterisation of perspicuity means that such judgments need not hold once and for all. Understanding, especially of the pragmatic kind championed by de Regt, is a context-dependent and agent-relative notion: a representation that enables us to better use the theory in one context need not do so in all contexts. And indeed, I argue, this is what we find in the examples discussed above. There are certain contexts in which the apparently less perspicuous representation better enables us to use the theory. In such contexts, it is that representation which most improves our understanding, and hence which should be regarded as the more perspicuous of the two. Finally, I argue that the highest level of understanding—the highest level of skill in theory-use—is contained not in any one representation, but rather in the capacity to be familiar with multiple representations, and to move swiftly and fluently between them. Hence, we arrive at something very like Weatherall's "puzzleball" account of explanation in physical theories (Weatherall 2017). Philosophers of physics, then, should not seek the one representation uniquely well-suited to depict the world: they should rather explore the menagerie of representations available, and seek to understand the relationships between them.

Jill North (Rutgers University)

On the Idea of Perspicuous Representations

There are in general different ways to represent something. Different mathematical devices can be used to represent the same mathematical or physical object. Different linguistic entities can represent the same content. Different mathematical formalisms can represent the same physical theory. More, the choice of how to represent something seems to be entirely under our conventional control. It seems we may stipulate that any representational vehicle be used to represent anything at all—as we can stipulate that a salt shaker is being used to represent Madagascar (Callender and Cohen, 2006). (Compare Putnam's (1983) thesis of trivial semantic conventionality and Teitel (2021).) There are also in general more or less good ways to represent something. The goodness of a representation might be entirely a matter of pragmatic virtues: some representations are more or less useful for creatures like us and the tasks we want to accomplish (as a geographical map is typically a more useful representation of Madagascar than a salt shaker is). Or the goodness of a representation might be more objective, flowing directly from the relationship between the intrinsic natures of the representational vehicle, on the one hand, and the represented item, on the other (as a geographical map in a certain sense better captures the nature of Madagascar than a salt shaker does; compare the idea of perspicuous representation in Møller-Nielsen (2017)). Or the goodness might arise from something in between—part pragmatic, part non-pragmatic, in any case not wholly objective. Philosophers have recently objected to the thought that there is any objective, intrinsic, and/or non-pragmatic sense of "perspicuous representation," for a variety of reasons. Barrett (2022) suggests there may be no suitably precise sense in which certain representations of mathematical structure are perspicuous and therefore privileged. Jacobs ((2022); this symposium) argues that perspicuous representations of spacetime structure require an extra ingredient (the dynamical laws), with the result that this is not an entirely pragmatic notion, but neither does it flow solely from the relationship between the intrinsic natures of the representational vehicle and representational target. Hunt (2021, 2022), in a different way, argues for something in between the pragmatic and non-pragmatic conceptions: a perspicuous representation can impart to us intellectual understanding, by making certain features manifest. Callender and Cohen (2006), Teitel (2021), and Dewar (this symposium) go further, arguing that there are no non-pragmatic constraints on the choice of representational device, and so no objective sense in which certain representations are more perspicuous than others. I defend the objective, intrinsic, non-pragmatic sense of perspicuous representations in the face of objections to the very idea of such a thing. Although this may require a level of realism that is anathema to

many philosophers of science, I aim to alleviate such qualms. I also link up the discussion to the topic of theoretical equivalence, via the question whether there is any interesting, non-pragmatic sense in which representations differing (only) in their perspecuity can reasonably be regarded as inequivalent.

Philosophically Analysing Expert-by-Experience Involvement in Psychiatry

Traditionally, psychiatric research has been conducted by psychiatrists on patients. This view has been heavily challenged by arguing experts-by-experience are an important source of scientific knowledge. Experts by experience should play a role in deciding what to research, how to research it and how to interpret the results. This raises many important issues related to philosophy of science such as about the nature of data, theory and methodologies. This symposium will use philosophy of science to analyse expert-by-experience involvement in psychiatric research, helping us understand this phenomena and allowing us to explore key philosophy of science topics in a novel domain. One paper considers how experts-by-experience relates to abstraction in science. Deciding what to abstract away cannot be established through the lived experience experts-by-experience provide. Also, lived experience is typically influenced by abstract science. One paper provides an alternative justification for the role of experts-by-experience. Expertise is something that emerges from a collective and dialogical pursuit, that involves a group reflection on evidence whereby group processes and facilitation support community consensus building. One paper considers how typical academic norms can downplay the role of emotionality in lived experience. This potentially blocks experts-by-experience generating valuable data and viewpoints which involve emotion.

Sam Fellowes (University of Lancaster)

How the Lived Experience of Experts-by-Experience Relates to the Abstract Nature of Science

In this paper I consider how the lived experience of experts-by-experience should relate to the abstract nature of science. I will argue that lived experience is actually influenced by abstract elements of science and argue that many aspects of constructing an abstract science cannot be derived from lived experience. One key justification of experts-by-experience being involved in psychiatric research is lived experience. An expert-by-experience has experiences which others lack. This means they can contribute unique data which other individuals cannot provide. However, significant parts of science involve abstracting from the data to create idealised models that do not accurately reflect the world. For example, physics includes models like frictionless planes, ideal gases and perfectly spherical objects which abstract away parts of the world. Psychiatric diagnoses are themselves idealised models that do not reflect particular people. They are generalisations that abstract away many aspects of particular people by not covering aspects of the individual like past life experience, present life situation and future life goals. A particular psychiatric diagnoses also abstractly join together aspects of different individuals since diagnoses typically cover many more symptoms than any particular individual with the diagnosis actually exhibits.

I argue that abstract scientific elements are already present when experts-by-experience draw upon lived experience. For example, if an autistic person gives a description of what it is like to be autistic we are, firstly, demarcating that person as being autistic rather than, say, schizophrenic and, secondly, the person is interpreting some experiences as being instances of autism. As such, abstract scientific elements are present when an expert-by-experience uses lived experience to provide data. Also, if we had an alternative diagnostic system which, for instance, had a single diagnosis that effectively joined autism and schizophrenia together then the individual would interpret their experience using different concepts. This means lived experience is not unmediated but is influenced by abstract aspects of science. This means the epistemological strength of lived experience is partly dependent upon the epistemological strength of the abstract science. I then consider what role experts-by-experience should play in formulating abstract science. Abstract scientific models are typically formulated by abstracting away parts of the data using general modelling principles. Experts-by-experience cannot directly use lived experience to decide which data to abstract away and how to model the remaining data. The data itself does not show which data needs abstracting away. I suggest that experts-by-experience should approve of typical modelling techniques for abstracting data which are typically employed within science and psychiatry. This significantly constrains which views of experts-by-experience should be considered scientific. In conclusion I argue that the data produced by experts-by-experience in psychiatric research actually only indirectly rests on lived experience. Lived experience typically needs interpreting through abstract scientific notions which are not derived from lived experience. Also, many important psychiatric notions are abstract models which cannot be derived from lived experience which potentially limits the roles of experts-by-experience.

Lisa Bortolotti (University of Birmingham), Michael Larkin (Aston University) & Michele Lim (University College London)

Expertise or Perspectives in Dialogue? The Role of Lived Experience in the Mental Health Context

Expertise is afforded - by people, questions, culture, expectations, power relations and systemic structures. These affordances can be changed - by skilful facilitation, preparation and expectation- setting and by thoughtful configuration of systems and structures - people's capacity for providing expert insight can be scaffolded or obstructed. If expertise is situated in this way, then we should choose to create

environments and relationships in which different perspectives can be heard and understood and the people providing those perspectives can be seen as experts. What makes someone an expert? Some form of relevant knowledge or experience is necessary for expertise but not sufficient for it. Expertise is built upon knowledge or experience, but it arises from extensive efforts in seeking to understand something; it is borne out of a 'hard-earned' familiarity with the contours of a particular set of problems. So, are experts-by-experience real experts? The objections raised to expertise-by-experience identify general worries about expertise in general, in terms of how to ensure objectivity, how to deal with disagreement, and whether the information or perspective provided is evidence for the claims the expert supports. Ultimately, one way to overcome some of these challenges to the authority of experts is to see expertise as something that emerges from a collective and dialogical pursuit that involves a group reflection on evidence. One way to bypass this constant requirement to justify the seat at the table for experiential insights is to talk about perspectives in dialogue rather than expertise. A perspective is a way of referring to how something appears from a particular standpoint, which acknowledges the relevance of that standpoint to what is foregrounded. A dialogue is a means of sharing insights, carried out to support reciprocal understanding. How do perspectives in dialogue work? Codesign and coproduction approaches encourage perspective taking and use group processes and facilitation to support community consensus building. Here we will offer an example of successful perspectives in dialogue from our own research.

Astrid Fly Oredsson (Independent)

Academic Affect Norms as a Barrier to Inclusion of Experts-by-Experience in Psychiatric Research

In this paper I argue that dominant affect norms in academic spaces can impede the participation of experts-by-experience and receptiveness to our insights. Specifically, since emotionality is frequently taken to be at odds with expert credibility, the 'difficult' testimony sometimes provided by experts-by-experience is likely to be seen as less credible than that of experts whose expertise is grounded in medical training, clinical practice and/or research. One notable methodological implication of this is that future philosophical work on the role of experts-by-experience in (psychiatric) research should investigate the interplay between affect norms and epistemic practices in academic spaces and research settings. A significant characteristic of academic spaces – seminars, conferences, lectures, journals and labs – is their stringent affect norms. These are not spaces suitable for strong displays of emotions. Rather, these spaces are pervaded by assumptions of so-called objectivity and professional distance between researchers and their research. Emotionality, on the other hand, is positioned as inconsistent with proper academic expertise, and emotionality is thought to undermine expert credibility. This, however, can lead to undue dismissals of valuable lived experience insights. While not all observations provided by experts-by-experiences are accompanied by high degrees of emotionality (in fact many of us are exceptionally skilled at speaking 'the academic language' e.g., using technical terms and providing testimony in accordance with dominant affect norms), some insight grounded in lived experience will be incredibly difficult to communicate without a certain degree of emotionality. Yet, in spaces where credibility is so closely tied to distance and 'objectivity', this type of emotionality is likely to be seen (consciously or not) as indicative of incredibility. Consequently, there is a viable risk that testimony provided by experts- by-experience will be assigned too little credibility merely because of the associated affect. This risk is particularly high when it comes to 'difficult' testimony pertaining to e.g., traumatic experiences. This phenomenon may result in a range of negative consequences (e.g., testimonial and affective smothering and marginalization, erosions in self-trust and hermeneutical injustice) all of which are likely to limit the participation of experts-by-experience in psychiatric research. Although a detailed examination of such consequences is beyond the scope of this paper, future philosophical work aimed at providing a thorough account of the in- or exclusion of experts-by-experience in research should pay more attention to the link between affective and epistemic norms. One obvious place to start would be to further analysis of the aforementioned potentially exclusionary upshots of dominant affect norms in academic spaces.

Large Scale Brain Models: From Technology to Biology

Large scale brain models simulate thousands of neurons, offering neuroscientists the possibility of seeing how cognition arises from their combined activity. This symposium brings together STS (science technology studies) and philosophy of neuroscience in order to examine the role of advanced computational technologies in the ongoing construction of theories of the neural basis of cognition. Tara Mahfoud presents results of ethnographic fieldwork in a computational neuroscience lab in which the boundaries between the brain and computer, biological and technological, are being negotiated in order to exploit the epistemic potential of large brain simulations. Mazviita Chirimuuta discusses the idea that the original explanatory targets of many natural scientific theories were technological objects. This sheds light on the role of deep artificial neural networks (ANN's) in neuroscience today. Rosa Cao discusses ways of accounting for similarities (and differences) between artificial neural networks and the biological brains that are the target of explanation. The role of ANN's in neuroscience today is controversial precisely because it is an open question how many of the neural details left out by these idealised models are relevant to cognition – a question that itself awaits progress in theoretical neuroscience.

Tara Mahfoud (University of Essex)

The Biological Imitation Game: Sublime Explorations of the Boundaries Between Human, Animal and Machine in Large-Scale Brain Modelling

Christoph Koch and Michael Buice from the Allen Institute for Brain Science published an article called "A Biological Imitation Game" (2015) in which they propose an alternative method of testing large scale brain models' validity. The goal of simulating the human brain in order to understand human behaviour and brain-related illnesses, they say, is an ill-defined goal because it is not clear what understanding entails, and what a complete model of a brain would be given incomplete knowledge of the human brain itself. They propose a variation on the Imitation Game developed by Alan Turing in the 1950s where the validity of brain simulations is assessed based on how long brain data from simulations can 'fool' experts into thinking it is data from real, living brains. This 'game' imagines a future where one will not be able to tell whether the results of an experiment were produced by recordings from an *in vivo* or an *in silico* microcircuit.

Based on ethnographic fieldwork in a computational neuroscience laboratory in Switzerland between 2014 and 2015, I argue that by paying attention to the aesthetics of brain simulation in the laboratory, to how neuroscientists visually assess the validity of their brain models, we get a sense of how the boundaries between human, animal and machine are disrupted in the very process of being made equivalent. I attend to moments where neuroscientists see the results of their simulations as equivalent to *in vivo* or *in vitro* experiments, and their assessments of how well they have reproduced the individual neurons they are reconstructing *in silico* and the behaviour of neural networks they are simulating. Many neuroscientists, philosophers and social scientists have argued that large scale brain modelling and other initiatives that are entangling cognitive science and computing with medicine and health is part of what has been called the 'project of convergence' or the 'technological singularity' – attempts to break down boundaries between the human and the machinic. However, my research suggests that neuroscientists involved in building the computational models of brains are consistently faced with the limits of their approach. Through sublime experiences, neuroscientists are trying to come to terms with the differences and similarities between their brains, the rodent brains they are studying, and the machine brains they are building. They are also trying to come to terms with the implications of their work – what does it mean to say they are building a simulation of a brain? In this interplay, the human brain is imagined and re-inscribed as a computational machine – but one that is different to the machines they are building in their lab.

Rosa Cao (Stanford University)

Measuring Similarity in Artificial Neural Networks and Biological Brains

As deep neural networks have become better at performing sophisticated tasks, they have been touted as promising models of the brain. Given the obvious and numerous differences between biological brains and these artificial models, these claims are contentious. It may be that the only real similarities between the systems are the ones that we have built in, and we cannot expect DNNs to shed new light on how biological brains work. And it might also be the case that even though similar behaviors are exhibited, they are produced in fundamentally different ways than they are in animals, once we look at how artificial systems actually work. How can we adjudicate questions about whether DNNs are similar enough to brains, in the right respects, such that they can serve as good explanatory models of human and animal cognitive capacities? One way to approach the problem is to construct a quantitative measure of similarity that can be used to assess how well the activities of a model can be used to predict the activities of a brain. Such a measure will force us to make explicit not only the phenomena that we are trying to explain, but also which features of a model we will hold responsible for explaining them (i.e. just the ones that will contribute to our measure). An example of this approach is the Brain-Score project (Schrimpf et al 2020), which advocates testing neural network models directly given a standardized set of behaviors agreed to (provisionally) characterize a capacity or domain of interest, and measuring how well they predict the neural data associated with those behaviors in animals. Going beyond Brain-Score, we would like a measure of similarity that is generally applicable, so that it can tell us (for example) how different visual areas in a mouse are from visual areas in a primate, or how visual cortex in a young brain might differ from that in an older one of the same species. Moreover, any measure of similarity should be validated on whether it reproduces verdicts that we have independent reason to believe (e.g. that human visual cortex is more similar to macaque than to mouse, or that V1 is more similar to V2 than to V4). It should treat models and their targets symmetrically – no differently than it would treat two animal brains. Perhaps no such measure can be found – if so, then we will have evidence for a fundamental difference between artificial DNNs and their natural targets. But if such a measure can be constructed: symmetric, generally applicable, and independently validated, then we can ask more directly: how good are DNNs as models of the brain? And we will have found a domain in which we can reason about artificial models and their natural targets on equal terms.

Mazviita Chirimuuta (University of Edinburgh)

'What I Have Not Made, I Do Not Understand': Explaining via Remaking the Brain

Grossmann's view was that a theory central to the development of modern science, classical mechanics, was originally a theory of the workings of artificial machines, subsequently applied to the world at large: 'in all these technological upheavals, [man, sic] acquired new, important material for observing and contemplating the actions of forces. In the machines, in the turning of the water wheels of a mill or of an iron mine,

in the movement of the arms of a bellows, in the lifting of the stamps of an iron works, we see the simplest mechanical operations, those simple quantitative relations between the homogeneous power of water-driven machines and their output, viz. those relations from which modern mechanics derived its basic concepts.' (1935/2009, 128). An advantage of beginning with the theorisation of artificial systems was that these presented the interaction of forces in a simpler and more readily comprehensible way than could be encountered elsewhere. This strategy was not confined to the physical sciences, and the use of machine models as proxies for organisms has been a mainstay of biological research in the modern era (Canguilhem 1965/2008). The brain is an organ of unfathomable complexity, perhaps making machine-models indispensable. This paper describes how a major strategy for neuroscientists confronted with the task of forming explanations of how the combined activity of populations of neurons enables cognitive performances is to use certain technological objects (artificial neural networks) as proxies for the brain itself. Though still very complicated, ANN's lack most of the features that complicate biological neural networks, such as vasculature, glial cells, and distinct types of neuronal morphology and physiology. It is argued that to a large extent, explanation in cognitive neuroscience is dependent on the ability of scientists and engineers to make artificial systems such as ANN's and other computational devices, which have some degree of functional similarity with the brain area under investigation. Understanding is secondary to, and dependent on, making. This strategy of explaining via (re)making will be related not only to recent pragmatist philosophies of science, but shown to connect in important ways with more radical views on human knowledge and understanding such as Vico's celebrated "verum factum" and Kant's proto-constructivism. I will discuss the implication that the science of complex neural systems is restricted in its ability to comprehend the brain "in itself", and offers explanations of the neural basis of cognition only through the lens of artifice.

The Many Faces of Empiricism: Physical Theories Between Idealism and Realism

The legacy of empiricism has been at the centre of many long-standing debates in the history of 20th century philosophy of science. However, the empiricist tradition has been the subject of diverse and disparate interpretations. Moreover, the notion of empiricism that is at the core of philosophical debate is often defined by the position with which it is contrasted, whether that be, e.g., critical idealism or scientific realism. The aim of this symposium is to explore the changing faces of empiricism in debates over neo-Kantian idealism, phenomenology, and scientific realism over the course of 20th century philosophy of science, up to the present. The symposium will focus on three case studies from physics, covering work from throughout this period, and based on these studies, we will look to provide a synoptic view of the central role of empiricism in the epistemology of science and promote dialogue between conflicting traditions.

Noah Stemeroff (University of Bonn)

Cassirer and Weyl on Helmholtz (and the Empirical Foundation of Scientific Thought)

Though representative of somewhat divergent intellectual traditions, Ernst Cassirer and Hermann Weyl both laid claim to aspects of the Helmholtzian epistemological program. Following Helmholtz's influential study of the problem of space (i.e. on the problem of which geometrical structure(s) can be taken to serve as the a priori ground for scientific cognition), both held that a given group-theoretic framework must always serve as a necessary presupposition of scientific thought (within a broadly Kantian tradition). Neither held that this framework was fixed, a priori, as it was for Kant. However, in allowing for the revision of the constitutive framework of scientific thought, both were forced to face the spectre of a pervasive relativism. In response to this challenge, Cassirer and Weyl both followed Helmholtz in suggesting that the relativist abolition of the 'absolute' standard of objectivity does not entail the abolition of the difference in value and performance of various scientific theories. Scientific theories do not stand apart in their relation to the 'world', to be judged solely on their own merits, but rather as part of a progressive series. In developing this aspect of Helmholtz's thought, Cassirer and Weyl both highlight the fundamental role that group theory plays as a constitutive feature of our understanding of objectivity through the progress of physical theory. In doing so, Weyl (in his later writings) sought to defend certain aspects of Helmholtz's 'Kantian' empiricism, by seeking out an empirical ground for the construction of a theoretical picture of reality. In contrast, Cassirer stripped the Helmholtzian program of its direct empirical grounding, on route towards a novel neo-Kantian structuralism. In this paper, I will provide an examination of the diverging views of Cassirer and Weyl concerning the constitutive role of group theory in physical enquiry, as a way of exploring their unique development of the Helmholtzian program and their particular understanding of the empirical foundation of scientific cognition. In particular, I will address the relevant notion of empiricism that is at the foundation of their theoretical programs (i.e. in both Weyl's and Cassirer's later thought, from the late 1920s onward). To conclude, I consider what lessons we can draw from this history concerning modern debates on empiricism, and its relation to critical idealism, scientific realism, and the methodology of modern physics.

Philipp Berghofer (University of Graz)

Realism in Quantum Mechanics? Lessons from Husserl's Empiricism

In my contribution, I distinguish two forms of empiricism, discussing what each implies for questions regarding the interpretation of quantum mechanics. First, there is what I

call Husserl's moderate empiricism. This empiricism is articulated in Husserl's Crisis (1936) when he criticizes Galileo's "mathematization of nature." This moderate empiricism, basically, consists of two claims:

ME1: Mathematical models and concepts are tools that allow us to represent aspects of reality, but we must not mistake this method for true reality.

ME2: The life-world, i.e., the world of our everyday experiences, is epistemically prior to the world of science.

ME1 implies that we must be careful when we reify or objectify mathematical concepts. ME2 puts the following constraint on our (interpretation of) scientific theories. If there is a clash between how scientific theories represent the world and how we experience the world, priority must be given to our experiences. Both ME1 and ME2 are highly relevant when it comes to contemporary philosophy of quantum mechanics. This can be best seen when discussing David Albert's 1996 version of wave function realism according to which the high-dimensional configuration space of quantum mechanics is the real space while our impression of living in three-dimensional space is "flatly illusory." Albert's approach perfectly exemplifies a popular strand in contemporary philosophy of science that has been captured by Sellars' saying that "science is the measure of all things," implying that the world of science is not only ontologically but also epistemically prior to the world of experience. It is in direct contradiction with ME1 and ME2. I take it that ME1 and ME2 constitute a moderate version of empiricism in the sense that most philosophers of science that identify as empiricists feel sympathetic to them. However, strongly influenced by Husserl's phenomenology, Hermann Weyl, for instance in his *Mind and Nature* (1934), has formulated an approach to science that embraces the following idea that I call strong empiricism.

SE: Scientific theories, at a fundamental level, do not represent objective reality but describe the relationship between the experiencing subject and the experienced world.

Weyl (and also Fritz London) believed that quantum mechanics can best be understood as exemplifying SE. Currently, the best developed and most consistent interpretation of quantum mechanics in this spirit is QBism. The distinctive idea of QBism is to apply a personalist Bayesian account of probability to quantum probabilities. This is to say that quantum states do not represent objective reality but instead represent an agent's subjective degrees of beliefs about her future experiences. Consequently, the quantum formalism is to be understood as a mathematical tool that allows the experiencing subject to predict what she will experience next. While there is some consensus that QBism constitutes a consistent interpretation of quantum mechanics, it is generally believed that it lacks a proper philosophical foundation. One objective of my contribution is to show how a phenomenological approach to science as articulated by Husserl and Weyl can help to make sense of QBism.

Guy Hetzroni (Open University of Israel)

Do Invariance Arguments Presuppose Realism?

The tension between realism and empiricism is one of the most debated issues in the philosophy of science, particularly in the context of contemporary physics. This paper takes a pragmatic- naturalist approach towards the issue, aiming to characterize realism in the context of theoretical physics by examining whether and to what extent do methods of theory construction in fundamental physics presuppose realism, and what is the relation of these presuppositions to formal mathematical notions on the one hand, and to empirical notions on the other hand. The empiricist-realist tension is to a great extent outlined by two notions: fallibilism and theoretical holism. The former is an empiricist desideratum that gives rise to the anti-realist argument of the pessimistic meta-induction. Theoretical and confirmation holism, associated with Duhem and Quine, is a thesis whose significance was greatly enhanced during the 20th century due to the growing dominance of mathematical reasoning. Quine argues for a theoretical holism that blurs the analytic-synthetic distinction, and further prevents reduction of statements to purely experiential facts. In this way this holism undermines the foundations of empiricism. Quine's argument, later developed by Putnam, uses the holism thesis as a part of an indispensability argument supporting mathematical realism. On the other hand, Duhemian holism greatly motivated forms of anti-realist approaches to science (e.g., Cassirer's), and is the basis for the anti-realist argument from underdetermination of theory by evidence. In this paper I will suggest an account of a central methodological thread in theoretical physics, through which the concept of invariance arguments was developed, first in the context of Einstein's special and general theories of relativity, and later in the context of abstract, internal symmetry groups in particle physics. This reflection on the methodology, it will be argued, shows that both fallibilism and holism are to a great extent inescapable. Theoretical holism does indeed characterize our theories, and fallibilism plays an internal role in the outlined theorizing. Yet, on the backdrop of this realization, I will show that these methods employ presuppositions that are best understood as manifesting a certain form of moderate and tenable realism that I would like to characterize based on its role in the practice of theorizing. This form of realism, I argue, is different from standard forms of scientific realism as well as from mathematical realism. While theoretical holism does characterize every given theory, methods of theory construction presuppose coarse separation of our "web of beliefs" into physical and mathematical parts. It is the mathematical part that is more subject to fallibilism. The physical part is based on non- mathematical notions such as contingency, measurability, locality and causality, and is conceived as the major guide to unobservable reality. This realism involves a notion of representation (of some aspects of unobserved physical reality by a mathematical structure) whose meaning is not defined *a priori*, but is rather constantly

examined as part of the scientific process. Empiricist desiderata play a crucial role in this progress, but they do not provide a sufficient basis for understanding it.

Major Transitions in Biology, Culture and Cognition

This symposium will analyse major transitions thinking across the domains of biology, culture and cognition. It leverages current work in two large-scale projects examining these issues: the Templeton-funded *Major Transitions in the Evolution of Cognition* project, headed by Marta Halina, Andrew Barron and Colin Klein; and the ERC-funded *Representing Evolution* project, headed by Samir Okasha. Three papers will be delivered, addressing the application of major transitions thinking in one of the three domains of biology, cognition and culture. Samir Okasha will present an overview of philosophical work on major transitions in evolutionary biology, and defend some answers to key questions in the literature. Andrew Barron will present a paper (co-authored with Marta Halina & Colin Klein) outlining an account of the major transitions in the evolution of cognition. Finally, early career scholars Arsham Nejad Kourki and Ross Pain will present an assessment of some key issues facing the application of major transitions frameworks in cultural evolution, and illustrate how interdisciplinary efforts can begin to overcome some of these issues.

Samir Okasha (University of Bristol)
The Philosophical Significance of the Major Transitions in Evolution

The contemporary interest in "major evolutionary transitions" (METs) can be traced to the pioneering works of Buss (1987), Maynard Smith and Szathmáry (1995), and Michod (1999). Though these authors' approaches differed considerably, both in respect of the empirical phenomena they were concerned with and in the type of explanation they sought, they converged on a number of key points. These included: (i) that a series of evolutionary transitions has occurred in the history of life on earth that radically re-shaped subsequent life forms; (ii) that some or all of these transitions involved formerly free-living entities coalescing into larger groups, giving rise to a new level of hierarchical organization; and (iii) that explaining how and why these transitions occurred represents an outstanding task for evolutionary biology. In the last twenty-five years many evolutionists have risen to the task, and the study of METs has flourished into a thriving research program, generating much interesting work, empirical and theoretical. While it is self-evident why METs are of scientific interest, it is perhaps less obvious why they should be of philosophical interest. And yet they are, as attested by the large body of literature on METs authored or co-authored by philosophers of science (Griesemer, 2000; Okasha, 2005; Godfrey-Smith, 2009; Calcott and Sterelny, 2011a; Birch, 2012, 2017; Clarke, 2014; O'Malley and Powell, 2016; Ryan et al., 2016; Currie, 2019). In this literature we can detect two distinct sorts of philosophical question: "conceptual" and "ontological". Examples of the former include questions about how exactly an evolutionary transition should be defined; what form an evolutionary explanation of an MET should take; how concepts such as kin selection, multi- level selection and the "gene's eye view" apply to METs; and whether an overarching theory of evolutionary transitions is possible. Examples of the latter include questions about what status the new biological units that arise from METs have (e.g., are they organisms, superorganisms, or individuals?); and about the nature of the hierarchical organization that results from the transitions (e.g., is it a hierarchy of parts and wholes? does it have a privileged level or are all levels of equal status?). Tackling these questions requires an integrative approach that draws on both biology and the philosophy of science. The aim of this talk is to provide an overview of these questions, to defend particular answers to some of them, and to illustrate by example how philosophical analysis can shed light on this important area of evolutionary biology.

Andrew B. Barron (Macquarie University), Marta Halina (University of Cambridge) & Colin Klein (Australian National University)
Major Transitions in the Evolution of Cognition

Recent findings in comparative cognition seem to have confused rather than clarified our understanding of the evolution of cognition. In a world of algebraic bees and smart slime moulds is there any pattern to the evolution of cognitive capacity? Maynard Smith and Szathmary (1995) famously provided a framework for understanding the grand scheme of biological evolution by positing a few major transitions - such as the origins of chromosomes and multicellular life – which enabled radically different forms of life to evolve. Here we propose that the evolution of cognition can also be comprehended as a series of evolutionary transitions. We argue each transition is a qualitative change in the computational architecture of a system. Selection for operational efficiency or robustness can drive structural changes in computational architecture that then make new types of cognition evolvable, and change what is possible for a lineage. We recognise five major transitions in the evolution of animal cognition. Each gives rise to a different type of computational architecture which changed the evolvability of a lineage and opened up phenotypic space to allow the evolution of new cognitive capacities. This transitional account helps us comprehend both the process of cognitive evolution in animals and the diversity of animal minds.

Arsham Nejad Kourki (University of Cambridge) & Ross Pain (Australian National University)
Major Transitions in Cultural Evolution: Key Challenges and New Directions

Since its original conception in the mid-90's (Maynard Smith and Szathmary 1995) the idea of major evolutionary transitions has undergone evolutionary trajectories of its own. The major transitions framework has diversified into multiple theories with varying degrees of divergence from each other and from Maynard Smith and

Szathmáry's original proposal. These range from the more mainstream theory of evolutionary transitions in individuality to various theories of major transitions in cultural evolution. But is the application of the major transitions framework to areas outside of biology warranted? In particular, the original framework was developed to explain processes of biological change; how then do we apply it to culture, given that processes of cultural change are very different to processes of biological change? In this talk, we will first defend a relaxed notion of major evolutionary transitions that specifies a minimal family of requirements for a transition to qualify as major. This move allows for the accommodation of a plurality of major transition frameworks both internal and external to biology. We then show how our conception maps on to a specific proposal from the recent literature; namely Birch and Heyes' (2021) account of the cultural evolution of cultural evolution. We show how conceptual work in the philosophy of biology applies to new work in cultural evolutionary theory, and how both can be supported by evidence, in this case from paleo-archaeology. We argue this demonstrates how major transitions thinking can be successfully applied beyond the domain of biology.

Similarity and Economic Models

This session will investigate a notion that is crucial to economic modeling, but one that is—in the economic-modeling context—philosophically under-explored: *similarity* between model and target. Robert Northcott explains why vague appeals to similarity are widespread in economics, but he argues that they are not sufficient. Benjamin Genta distinguishes between two types of analogical inference, argues that philosophers have focused on only the first of these types, and shows that the second type is an important way that economists reason from models. Igor Douven presents a formalization of reasoning with similarity, and he uses this formalization to question the normative status of analogical reasoning. He then applies this analysis to a publicly available dataset. Nadia Ruiz argues that both a similarity-based approach of model evaluation and its adequacy- for-purpose rival have important limitations when it comes to model-building in economics. This is because the data that economists have at hand is often incomplete.

Igor Douven (Université Paris 1 Panthéon-Sorbonne - CNRS)
Analogical Reasoning: A Carnapian Approach

Analogical reasoning exploits similarity relations among items in a domain of interest or among properties the items may possess. For example, from the fact that your friend likes oranges, you may be inclined to infer that she will like tangerines, too, given that tangerines do not taste that differently from oranges. Students of analogical reasoning face two tasks: (i) making precise the role that similarity plays in inferences of this type, in particular clarifying the connection between similarity and strength of inference, and (ii) charting the norms that such inferences are to obey if they are to qualify as valid or rational (your expectation that your friend will like tangerines appears reasonable, which indicates the presence of a normative dimension). There is still no better starting point for tackling these questions than Carnap's late work on inductive logic. In this work, Carnap abandoned his earlier project of defining inductive logic in strictly syntactical terms and instead tried to exploit semantic relations among predicates. For this purpose, he appealed to (what he called) attribute spaces, which are mathematical spaces meant to represent similarity relations among items and properties, hoping that these would help formalize analogical reasoning. Attribute spaces remain rather abstract and underdeveloped in Carnap's work, but they can be thought of as the precursors of what are now commonly known as conceptual spaces. In psychology, a number of such spaces have been studied in detail over the past twenty years, and we also possess a deep theoretical understanding of conceptual spaces. In my talk, I look at the prospects of pursuing a Carnapian approach to the study of analogical reasoning within the conceptual spaces framework. Drawing on insights from the conceptual spaces literature, I intend to show how we can formalize analogical reasoning in a way that is both mathematically precise and psychologically plausible, thus contributing to the first of the above tasks. As for the normative status of analogical reasoning, various authors have rightly pointed out that a proper formulation of the norms for analogical reasoning is still missing from the literature. Using computer simulations, it is shown why this is unsurprising, in particular, that the normative status of analogical inference is domain-dependent so that one cannot expect universally applicable norms of analogical reasoning. Finally, I take a first look at the descriptive adequacy of the new version of the Carnapian approach by working out some of the empirical consequences of the proposal and testing them by reanalyzing a publicly available dataset.

Benjamin Genta (University of California, Irvine)
Inferring Relations by Analogy

Analogies and analogical inferences play a vital epistemic role in both everyday and scientific life. The main contribution of this presentation will be in carefully distinguishing two types of analogical inference: outcome- and relational- analogical inferences. These types, I show, differ in *what* is being inferred. Philosophers [Bartha 2010; Hesse 1966] and cognitive scientists [Gentner 1983] have focused their analysis on outcome inferences. As I will show, however, relational analogical inferences are an important tool of scientific practice—especially in economics and other model-based disciplines. I will use Schelling's [1969, 1971] segregation model as a case-study. The talk is structured as follows: first, I give an account of three necessary steps needed to make an analogical inference. These are: (i) the representation stage, which is the process of formalizing the source domain and the target domain such that they are in comparable languages; (ii) the analogy mapping stage, where we draw the relevant

similarity matches between the source and the target; and, (iii) the analogical reasoning stage, where based on the similarity map drawn in the previous stage, we reason about the target domain. This reasoning can, for example, either be an inference or an explanation. These steps are based on previous work on analogy but have not been organized in this stepwise manner. Then, I present a standard account of analogical inferences—these are the outcome analogical inferences. In particular, these are inferences that infer an outcome in stage (iii). I then move on to give an account of a distinct form of analogical inference, which I call relational analogical inferences. One way to distinguish relational-inferences from outcome-inferences are that relational-inferences are inferring part of a causal mechanism rather than a particular outcome. In other words, it is not a concrete fact about the world that is being inferred, but rather an abstract relation between two existing entities. Prominent accounts of analogical reasoning seem to forbid such an inference, but as I show—using Schelling's [1969; 1971] segregation model as a case study—these relational inferences are widely used in model-based disciplines such as economics. Before concluding, I further distinguish some of the unique features between outcome- and relational- analogical inferences. In particular, they differ by (1) the background principle that warrants the inference, and (2) the types of investigation that these inferences are used for.

Robert Northcott (Birkbeck, University of London)

Similarity and Fragility

Similarity between model and target is a central concern of economic modeling. Not just any kind of similarity will do. There are two distinct dangers. I argue that each of them stems from the same methodological mistake. 1. Vagueness. A claim of similarity – and thus of explanation – must involve more than drawing a vague and intuitively appealing analogy between model and target. Yet such vagueness is widespread. But without empirical confirmation to back it up, we risk being led badly astray – and often are. 2. Fragility. Many, even most, relations in economic systems are fragile – that is, roughly, they hold only intermittently and unreliably. In these cases, models can be similar to their real-world targets at best only intermittently and unreliably. Therefore, explanations based on those models cannot just be assumed to be true generally because of empirical success on one occasion. They require fresh empirical warrant each time.
If relations in a domain are stable, then developing a master model is an effective methodological strategy. Warrant from empirical confirmation in one case will hold good across other cases, because of stability. A model may then be developed in relative isolation from continuous empirical refinement, confident that the stable relations it models hold reliably. If relations are fragile, however, matters are different. No model applies reliably, and emphasis must shift to contextual investigations, establishing each time which from a wide range of models might apply. Continuous empirical refinement becomes crucial. Developing a master model in isolation from empirical refinement risks that model becoming an irrelevant house of cards. The underlying danger is that a "master-model" strategy is mistakenly pursued even when relations are fragile. This is the common root of the other two dangers. First, it leads to vagueness because the empirical link between model and target becomes frayed; and second, it leaves a model vulnerable to fragility. Further, we see why the similarity issue is especially germane to economics – because economic systems often feature fragile relations yet are mistakenly tackled using a master-model methodology. The story carries a prescriptive punch: accepting claims of "similarity" without continuous empirical confirmation encourages an unproductive modeling strategy. So, we should not accept such claims.

Nadia Ruiz (Stanford University)

Beyond the Adequacy-For-Purpose View: Model-Building Constraints

The similarity view - in which similarity between a target-system and model does not ground representation but is a standard of accuracy [Frigg & Nguyen 2017] - has been overshadowed by the Adequacy-For-Purpose View (AFPV). The latter account, in contrast, focuses on the relationship between "target Y, user U, methodology M, circumstances B, and goal P jointly" [Parker 2020, 462] to evaluate the model's fitness for a specific purpose. Regardless of the differences between these two model assessment accounts, it is crucial to establish models' relation with real-world phenomena to assess scientists' communicative obligations in policymaking. This paper assesses the limits of current accounts of scientific assessment such as AFPV and the similarity view, by showing that neither account considers the constraints scientists face throughout model construction. These accounts have put too much weight on *how* values/purposes shape scientific practices, so that constraints such as measurement issues with index numbers, limited mathematical structures, incomplete and preliminary datasets, etc., have been neglected. This is problematic because these constraints sometimes, I argue, restrict scientists' or policymakers' purposes. For example, preliminary and incomplete data affect how economists analyze models that are used to make a policy change, specifically monetary policy [Diebold & Rudebusch 1991; Croushore & Stark 2000]. If only a preliminary incomplete dataset is available, how should macroeconomists assess whether this dataset is representative enough for the purpose at hand? In this case, limited data is a model- building constraint that does not allow the model's purpose- either epistemic or non- epistemic - to get off the ground, and it makes the model *unsuitable* for a model assessment such as AFPV. Focusing on the limitations posed by available datasets in macroeconomic modelling- cases in which the available data used to test the model is not the final or revised dataset- I illustrate the need to incorporate model-building constraints when assessing scientific practices. To discuss this is important because the possible constraints of preliminary data - my focus in this paper although it is not the only possible constraint when model-building - entail that neither the similarity view nor AFPV can assess models' fit adequately. In

other words, these accounts do not give us an answer to questions such as: what model must economists use to inform policy? Model-building constraints, in some cases, restrict scientists' and policymakers' epistemic and non-epistemic values/purposes, telling against any accuracy-for-purpose model evaluation account. Since model-building constraints are often present, giving up on any model evaluation account is not the solution. Instead, economists (scientists), besides evaluating models (scientific practices) should discuss explicitly how such constraints bound models' epistemic import.

Quantitative Approaches to Science at Scale

Scientists and administrators have turned to a variety of quantitative tools to deal with the deluge of research that is characteristic of modern-day science. In this symposium, we critically evaluate several such tools, with a particular focus on new and potential strategies for managing grant proposal review and addressing the replication crisis. Creel and Bright argue that automated grant review using machine learning would reduce beneficial epistemic diversity in science. Mayo-Wilson examines the use of across-reviewer averaging in proposal review and proposes other social choice procedures that would work equally well. Dang and Ward discuss the epistemic value and significance of multi-analyst studies, which crowdsource statistical analysis among dozens of research teams. The symposium aims to enrich our understanding of the social epistemology of science and advance efforts to improve its organization.

Kathleen Creel (Northeastern University) and Liam Kofi Bright (London School of Economics)

Don't Use Machine Learning to Evaluate Grants

Funding science is a chancy business. Promising projects come to naught; strong results fail to replicate. To reduce the uncertainty of their bets, grant-making agencies such as the Defense Advanced Research Projects Agency (DARPA) are encouraging the development of tools that aim to predict which papers will replicate - and, eventually, it is hoped, which grants DARPA should fund. The first generation of such tools use machine learning to assign each paper or grant a reliability score (Gordon et al., 2020; Yang et al., 2020; Alipourfard et al., 2021). This score indicates the likelihood that the result is replicable and robust. A lab with many papers judged as replicable will be more likely to be funded. The benefits of a tool that could predict the success of funding proposals seem clear: time saved, public funds better allocated. But in this paper, we argue that using machine learning to construct such a tool poses a significant risk to the epistemic diversity of science. As has been well argued, scientific inquiry functions best when it covers logical space, including exploring possibilities that may be or seem antecedently unlikely (Dang and Bright, 2021, §5). Failing to do so risks missing out on fruitful explanatory theories (Stanford, 2006; Longino, 2013). Encouraging a variety of scientific approaches to inquiry also insures us against error (Solomon, 1992; Zollman, 2010). And finally, ensuring a variety of approaches within the scientific community has been argued to be necessary for scientific objectivity (Neurath, 1946; Longino, 1990; Harding, 1995; Weatherall and O'Connor, 2020), evidential support, and decision-making under uncertainty (Lloyd, 2010; Dang, 2019; Heesen et al., 2019; Roussos et al., 2021). Why will filtering grants with machine learning reduce the diversity of scientific approaches? First, machine learning is data hungry. If every algorithm is trained on the corpus of accessible scientific papers, then even different algorithms trained by different research groups are likely to produce similar outcomes, judging the same papers to be unreliable (Bommasani et al., 2022). These correlated outcomes will result in correlated errors: different funding agencies relying on machine learning to screen grants will each erroneously reject the same proposals. To the extent that reliance on machine learning homogenizes outcomes, increasing the consistency of rejection for some projects and acceptance for others across repeated attempts to gain funding, it will decrease the epistemic diversity of science compared to the status quo. We argue that this disbenefit is not worth even significant increases in predictive success, proposing instead funding allocation strategies that benefit from predictive tools while maintaining a diversity of scientific approaches.

Conor Mayo-Wilson (University of Washington, Seattle)

Peer Review and Social Choice

Grant-funding agencies like the National Institutes of Health (NIH) exert enormous control on scientific research agendas. Moreover, because policy-makers face increasing pressure to rely on sound science, grant-funding agencies also exert substantial control on public policy. Philosophers of science, therefore, have an important role to play in (i) articulating the democratic and epistemic aims of grant-funding mechanisms and (ii) comparing *existing* funding mechanisms with *possible* alternatives that might achieve those aims better. Existing funding mechanisms are often deceptively simple. Individual reviewers first score and write brief reviews of proposals. To *amalgamate* those reviews, funding agencies like the NIH often *average* reviewers' scores and then make a holistic decision that is largely informed by those averages. What, if anything, justifies averaging, and can funding agencies do better? Drawing on interdisciplinary work in statistics, economics, and political philosophy (e.g., see Harsanyi (1955) and Muliere and Parmigiani (1993)), we first provide a defense of averaging from a social choice perspective – characterizing the "axioms" of social choice (e.g., neutrality, independence of irrelevant alternatives, etc.) that are necessary and sufficient for justifying a broad class of averaging functions. We then *generalize* other common social choice rules (e.g. Copeland's method) so that they can be applied in grant proposal review, and we *compare* the virtues and drawbacks of those rules to averaging by assessing which "axioms" each social choice rule satisfies. Our work contributes to the

rapidly growing body of philosophical work on peer review (see for example, Lee et. al. (2013), Lee (2015), and Heesen and Bright (2021)). Generalizing and evaluating existing social choice rules raises at least two philosophical difficulties. First, social choice theorists typically assume that all voters rank or score *the same set of options/candidates*. Because of conflicts of interest and differing areas of expertise, however, different reviewers often assess different grant proposals. Thus, to compare averaging to other social choice rules, one must first state plausible axioms for social choice in settings in which reviewers are *prohibited* from evaluating some options. This task is especially difficult because a fundamental aim of many democratic procedures is to treat voters "equally" (and this goal is often formalized in an axiom called "anonymity"). In grant review, however, voters/reviewers are not equal. Second, axioms for social choice rules are typically construed as constraints on *fair* democratic decision making (see Pildes and Anderson (1990)), not as properties of *epistemically reliable* procedures. Yet grant-funding agencies desire to judge chances of success of various research projects, not just whether the goals of such projects advance democratically endorsed goals. Although there are some epistemic justifications for existing social choice rules (List and Goodin (2001)), such justifications are typically *asymptotic*: they assume the number of voters is very large, as is the case in national elections. Grant proposal review, by contrast, typically involves only a handful of reviewers. We conclude, therefore, by discussing some alternative ways of evaluating the epistemic quality of various funding mechanisms.

Haixin Dang (University of Nebraska Omaha) & Zina Ward (Florida State)
Multi-Analyst Studies and Permissive Evidence

Are football referees more likely to give red cards to players with dark skin than to players with light skin? This was the question a pair of researchers sent to 29 research teams in 2014, along with a large dataset that included information from four major football leagues about referee calls, player demographics, and how often referees encountered each player, as well as a rating of each player's skin color. The research teams were asked to statistically analyze the data and determine whether skin color affects red card calls. 20 of the 29 teams found a statistically significant effect; as Silberzahn and Uhlmann (2015) explain, "findings varied enormously, from a slight (and non-significant) tendency for referees to give more red cards to light-skinned players to a strong trend of giving more red cards to dark-skinned players" (190). This sort of variation is often a feature of such "multi-analyst studies," which have spread to fields including neuroimaging (Botvinik-Nezer et al. 2020), biomedicine (Bastiaansen et al. 2020), psychology (Schweisenberg et al. 2022, Hoogeveen et al. 2022), and social science (Breznau et al. 2022). In all of these studies, a large number of research teams are given data and asked to assess one or more hypotheses. Most such studies return a surprising variety of answers. In this paper, we first survey deflationary explanations for why analysts reach different conclusions from the same data: the data might be noisy or indeterminate; the effect size might be extremely small; the target hypothesis might be poorly formulated; the data may be not probative of the target hypothesis; and some analysts may simply be making errors. While each of these factors plays a role in some cases, we suggest they cannot explain away all of the variation in multi-analyst studies. For instance, there can be multiple conflicting, yet equally valid, ways of operationalizing the same variable, leading to an ineliminable element of choice in determining the evidential relevance of a dataset (e.g., Schweisenberg et al. 2021). This "leftover" variance component, we suggest, is best explained by the view that evidence is permissive: it need not dictate a single rational doxastic attitude. In epistemology, this view is known as "permissivism" (White 2005). Our argument constitutes a novel defense of permissivism, one grounded in contemporary scientific practices and the ineliminable "subjectivity" in statistical inference (Rouder et al. 2016). It is tempting to see permissivism as the first step on a slippery slope to epistemic relativism. We argue, however, that multi-analyst studies also illuminate a valuable strategy for managing the uncertainty that arises in determining the evidential relevance of data. These studies probe the limits of what a given dataset can support, improving our understanding of how small methodological choices lead to different conclusions. We claim that this defuses some objections to permissivism in the philosophical literature. Although evidence is permissive, multi-analyst studies can help establish the boundaries that constrain rational believers.

Arrows of Time

The symposium discusses different ways of defining the direction of time through grounding in physical quantities, i.e., different definitions of Arrows of Time. Bryan Roberts will discuss whether reversibility is a problem for the reduction of thermodynamics to statistical mechanics. Katie Robertson will give an account of 'simple' initial states in Gibbsian statistical mechanics. Dominic Ryde will present an analysis of the proposal to define the Arrow of Time within the scale-invariant model of Newtonian mechanics. Lena Zuchowski will present an analysis of how important notions of randomness and disorder are for definitions of the Arrow of Time.

Bryan Roberts (London School of Economics)
Is reversibility a problem for the reduction of thermodynamics to statistical mechanics?

The 'reversibility problem' for reduction states that the laws of thermodynamics seem to be time-asymmetric in situations where the laws of statistical mechanics seem to not be, unless further structure like special initial conditions are added to the latter. I contend that on closer examination, there is no qualitative difference between the temporal symmetries of thermodynamics and of statistical mechanics---not even in the

absence of initial conditions. Thermodynamic time asymmetry emerges from an incomplete description of reality, which physicists call an "open system", for which statistical mechanics is time-asymmetric too.

Katie Robertson (University of Birmingham)
Discovering Simplicity
Accounts of time-asymmetry often require an initial condition. The nature of this initial condition varies: for Boltzmann and Albert, it is a low entropy initial state of the universe, for Wallace, the initial state must be 'Simple'. But more needs to be said about how to specify simplicity and discovering which states are simple. In this talk, I draw on resources from the causal modelling literature and connection to algorithmic complexity highlighted by Williams (2022) to give an account of 'simple' states in Gibbsian statistical mechanics.

Dominic Ryder (London School of Economics)
Directed Temporal Asymmetry from Scale Invariant Dynamics: Is the Problem of Time's Arrow Solved?
The scale invariant model of Newtonian gravity by Barbour, Koslowski, and Mercati purports to solve the problem of the arrow of time but has received minimal philosophical analysis. This omission is amended in the present work, in which I describe how the model manages to derive asymmetric behaviour from symmetric physics. The Janus point structure of the proposed solution holds significant preliminary promise for deriving asymmetry from symmetry. However, the proposal does not recover sufficient supervenience relationships between various other arrows of time to regard the problem as being solved. I propose a line of research which defines statistical mechanics within the BKM model, which would be a significant advance toward a satisfactory solution.

Lena Zuchowski (University of Bristol)
From Randomness to the Arrow of Time
This talk will demonstrate the importance of the historical and conceptual foundations of statistical mechanics by mapping out different derivational routes to the Arrow of Time. It will demonstrate that there are at least three different derivational routes: (i) starting from the Thermodynamic Entropy and inductively deriving the Empirical 2nd Law to ground the Empirical Arrow of Time; (ii) starting from a notion of randomness, which acts as a desiderata on the definition of Boltzmann and Gibbs Entropy, from which one can deductively derive the Statistical 2nd Law and then ground either (ii) the Universal Statistical Arrow of Time, or (iii) the Local Statistical Arrow of Time. Each of the three Arrows has different epistemic advantages and disadvantages: prominently, the Empirical Arrow of Time provides a straightforward definitional grounding of the direction of time; in contrast, the two statistical Arrows of Time have higher explanatory potentials, but their derivation requires the introduction of additional assumptions.