

GWP.2022

Abstracts of Contributed Papers and Symposium Contributions

- alphabetical by contributor -

A

Ariew, André – see Desmond, Hugh

B

Balorda, Vito & Šustar, Predrag: “Natural Selection: Pathway or Mechanism? Insights from Cancer Research“

In this paper, within the debate on the nature of natural selection, we examine causal approaches, that is, the new mechanistic accounts and the pathway concept. The mechanistic account has advanced both positive and negative assessments whether the mechanism concepts ‘get at’ the nature of natural selection. However, apart from only a few accounts that characterize natural selection along the lines of the main mechanism concepts, other accounts have mostly been cautious or, even, skeptical about such full-fledged mechanistic characterizations (for the skeptics, see, e.g., Havstad 2011; Garson 2021). Thus, contrary to some other areas of biological research, molecular biology to a highest degree, evolutionary biology appears to be out of a direct mechanistic reach. We examine the abovementioned negative assessments and propose a new perspective on the nature of natural selection by considering a current resurgence of interest in causal pathways. The causal pathway concept, as described by Ross (2021: 137): “refers to a sequence of causal steps that string together an upstream cause to a set of causal intermediates to some downstream outcome” (e.g., gene expression pathways, cell-signaling pathways, metabolic pathways, ecological pathways, and developmental pathways). Interestingly, Ross’ corresponding account does not address the causal pathway concept in evolutionary biology, meanwhile, as noted above, elaborates on all the other major biological disciplines. In the present paper, we examine the applicability of the causal pathway concept to the way in which natural selection is at work. In addition, we confront that causal pattern to the mechanism patterns that have been advanced for the nature of natural selection. We further test the applicability under consideration by taking into account the research area of ‘cancer as a micro-evolutionary process’. In other words, selection is present in

carcinogenesis once mutations to cancer cells (or proto-cancer cells) are heritable and make a difference to the relative success of cells, more specifically, cell lineages (see Plutynski 2018: 167). Despite some disputes about the extent of dissimilarities between that area and evolutionary considerations referring to individual organisms in a more standard sense, we argue in favor of the suitability of this case study in ‘getting at’ the nature of natural selection for the following reasons: (1) the accessibility of evidence with regard to natural selection, in this case, the population dynamics of independent cancer cells and their cell lineages; and (2) a specific coexistence of the mechanistic, pathway and other arrangements, as shown by the research area in question.

Baraghith, Karim & Feldbacher-Escamilla, Christian: “From Reduction to Unification: The Case of Cultural Evolutionary Psychology”

Cultural evolutionary psychology (Heyes 2018) accounts for the cultural evolution of cognition. It is based on evolutionary psychology and cultural evolutionary theory and aims at unifying both in a synthetic attempt. In this paper, we will show that, in sharp contrast to the reductionism of classical evolutionary psychology, cultural evolutionary psychology provides a unification. As we will argue, the form of its unification is ‘evidential’, and this form is to be preferred against purely ‘structural’ unifications as performed by competing approaches such as ‘dual inheritance theory’ in the nature-culture domain. The main difference between evidential unification and structural unification is that the latter ‘merely’ creates an abstract overarching framework for hypotheses and theories under consideration, without establishing a dependence relation between the respective kinds of evidence. Evidential unification, however, establishes a (mutual) dependence relation between different kinds of evidence and by this brings in further explanatory power.

Bartels, Andreas: “Kinematical constraints: No support for non-absolutism about laws of nature” (Symposium “Are all Laws of Physics Created Equal”)

Recently, Hirèche et al. (2021b) have proposed a new way of drawing the metaphysically necessary-contingent-distinction for laws of physics, which is based on the distinction between kinematical and dynamical structure. Referring to Curiel (2016) they argue that the crucial kinematical structure of physical theories consists in kinematical constraints – a main example for Quantum Mechanics is Pauli’s exclusion principle (PEP). I agree that kinematical constraints can be interpreted as expressing essential properties of the very systems a physical theory is designed to represent, and that by virtue of featuring essential properties of those systems the status of metaphysically necessary propositions can be assigned to them (even if this ontological interpretation presupposes some idealization with respect to a ‘final theory’). But, as I will argue, kinematical constraints, despite being candidates for metaphysical necessity, cannot provide support of a non-absolutist theory of laws. The reason for this is that kinematical constraints are not laws. As I have argued in Bartels (2019), the laws of physics are dynamical laws specifying kinds of processes (e.g. by characteristic Lagrange-functions). Possible counterexamples (‘constraint laws’, ‘composition laws’ and symmetries with their derived ‘conservation laws’) can either be classified as dependent parts of the mathematical apparatus of dynamical laws, or, in the case of symmetries, as representing general properties of physical laws, which means that they do not determine any specific dynamics. In particular, since PEP is derived from a symmetry requirement (permutation symmetry for many particles quantum states) characterizing kinds of states to be addressed by the theory (e.g. fermions), but not determining specific dynamics, PEP classifies as a kinematical constraint, but not as a law of physics. Therefore, non-absolutism about laws cannot be based on the kinematical/dynamical distinction.

Blanchard, Thomas & Hüttemann, Andreas: “Causal Bayes Nets, Causal Exclusion, and Symmetric Dependence”

Gebharter (2017) has argued that the causal Bayes nets (CBN) framework vindicates Kim’s exclusion argument. Consider the following Kim-style diagram:



There M_1 and M_2 are multiply realizable properties that happen to be realized by physical properties P_1 and P_2 (respectively) on the relevant occasion. (Black arrows represent metaphysical dependence, and the blue arrow represents causation.) Gebharter’s argument is that adding a causal arrow from M_1 to M_2 would yield a graph that fails the Minimality condition, and is thus prohibited by the CBN framework. Our goal in this paper is twofold. First, we show that if this argument is correct, another, larger exclusion problem looms: the CBN framework entails that *wholes* are causally excluded by their parts. In a nutshell, this is because a correlation between the behavior of a composite object and a putative effect of that object can be fully accounted for by positing arrows from the object’s parts to the effect; adding an arrow from the composite object itself to the effect is superfluous, and hence prohibited by Minimality. Thus, the CBN framework yields not one, but two exclusion problems. Our second goal is to offer a unified solution to these exclusion problems. Starting with parts and wholes, we think the solution lies in noting that the problem only arises on the supposition that wholes *asymmetrically* depend on their parts. Within the CBN framework this assumption is debatable. The CBN account is intended to be a lean metaphysical framework closely tailored to scientific practice. But arguably, a minimal metaphysics need only posit a symmetric dependence relationship between parts and wholes to make sense of scientific practice concerning parts and wholes (Hüttemann 2021). Moreover, on the standard interventionist interpretation of CBNs, causal Bayes nets have the function of capturing manipulability relationships. Because part-whole relationships are manipulable in both directions, this supports modeling them as symmetric in CBNs. Once this is done, we claim, CBNs fully vindicate the causal efficacy of composites. More precisely, once we extend the standard axioms of CBNs (the Markov and minimality conditions) to graphs containing such symmetric dependence relationships, we find that the relevant axioms vindicate the causal efficacy of composites, or so we will argue. In the last part of the paper, we show that this account also provides a solution to the exclusion problem for multiply realized properties. Specifically (we claim), the arguments in favor of regarding the dependence between wholes and their parts as symmetric also apply to the relationship between multiply realized properties and their realizers. The upshot is that in Kim’s scenario, we should represent the relationship between P_1 and M_1 (and between P_2 and M_2) with an undirected edge, not with an arrow. Applying our extended CBN framework to the resulting graph yields the result that M_1 is a cause of M_2 , so that the causal efficacy of multiply realized properties is vindicated. We close by indicating how our account compares to Eva and Stern’s (2020) solution to Gebharter’s argument.

Blečić, Martina & Šustar, Predrag: “Biological Metaphors as Vehicles for Explanation?”

It is generally claimed that a term or a specialized linguistic unit differs from a common language word by its unambiguous relationship with the concept it signifies and the stability of the relationship between form and content, but this is an idealized vision of specialized communication. Science is riddled with metaphors, which presents a problem for this view of scientific language. Now, what role do metaphors play in science and why are they maintained in such a specialized domain? The answers usually emphasize the following potential theoretical roles for scientific metaphors: (1) description; (2)

explanation; (3) prediction; and (4) a heuristic role (see Stegmann 2016). Additionally, Camp (2020) proposes a more general, (5) framing role. Here, we will focus on (2) and (5), in particular as they relate to our understanding of some basic biological phenomena. – Stegmann (2016) defends an explanatory account of scientific metaphors, esp. the “code” metaphor in molecular biology. According to him, coding schemes provide mechanism sketches that can have an explanatory role. We depart from that account for the following reasons: (1) it is not clear how the metaphor in question, as well as other metaphors (e.g., the genetic information metaphor) are related to explanatory structures; and (2) “code” and other cognate metaphors are replaceable with other, semantically less loaded, notions such as causation. – Levy (2020) endorses Stegmann’s idea and claims that metaphorical descriptions can be explanatory to the extent that they succeed in enhancing understanding. Following Camp, he claims that metaphors frame a target and thereby enhance our ability to think about it. According to him, this is how they explain. – We propose a reading of Camp that is not in accordance with this conclusion. Camp argues that metaphors play a fruitful role in science because they are intuitive and only partially consistent; they engage imagination, guide attention, and suggest hypotheses. These features make them useful scientific tools, but simultaneously distinguish them from explanations. They are framing devices, i.e., representational tropes that guide our interpretation by providing a perspective. Thus, metaphors can play a useful epistemic role in science, and can eventually lead to explanations, but cannot be equated with them. – Accordingly, we argue that Levy deflates the notion of explanation and that Stegmann draws a connection between metaphors and explanation which is too direct. We will try to show that indeed metaphors can play an important conceptual role in science and understanding but that their relation to explanation is not so close as the accounts considered here propose.

Blessenohl, Simon & Sarikaya, Deniz: “A Norm for Science Advice: Making Beliefs Accurate”

How should scientists communicate their findings when they advise politicians? One view holds that scientists should say what they have a high credence in. For example, they should not assert ‘X is toxic’ if they only have a credence of 0.7 that X is toxic. Rather, they should make their uncertainty explicit to say something weaker that they have a high credence in, such as ‘it is likely that X is toxic’. Another view holds that scientists should say what they expect to have the best policy consequences. For example, if scientists know that politicians will enact climate policies only if scientists do not make their uncertainty explicit, and the scientists think that climate policies are desirable, then they should not make their uncertainty explicit. We explore a third view, according to which scientists should say what they expect to make the politicians’ credences most accurate. That is, if a scientist has a credence of 0.7 in X being toxic, then she should say whatever brings the politician’s credence close to 0.7. If this requires not making uncertainty explicit or saying things she takes to bring about suboptimal policy consequences, so be it. For ease of reference, let us state the three views as three norms for scientists advising politicians.

(Honesty) Advising scientists ought to say what they have a high credence in.

(Policy) Advising scientists ought to say what maximizes the expected value of the policy consequences of what they say.

(Addressee) Advising scientists ought to say what maximizes the expected accuracy of their addressee’s credences in the target propositions.

The three views are simplified versions of more plausible views. This talk explores the advantages, problems, and implications of (Addressee). First, we outline potential advantages of (Addressee) over its two alternatives. In particular, (Addressee) does not permit scientists to skew their advice based on their moral assessment of policies. This is an advantage over (Policy) because such influence would undermine procedural values of democratic decision-making. It also constitutes a defense of the value-

free ideal, which is sometimes attacked on the basis of a norm such as (Policy). Also, (Addressee) does not require scientists to say what they have a high credence in even if that is counterproductive to induce more accurate credences in the addressee. This is an advantage over (Honesty) because, in such cases, it seems permissible to say what one has a low credence in, if that makes the politician's credences more accurate. We then turn to a central problem of (Addressee), that it seems to require vicious communication strategies if those happen to maximize the expected accuracy of the politician's credences. Finally, we assess prominent examples of science advice in the light of (Addressee). We are interested in what should guide scientists when they communicate to politician's in standard, non-pathological situations: honesty, policy, or accuracy.

Borner, Jan: “Causal Power Quantified - A Generalisation and Defense of Cheng’s Causal Power Measure”

As part of her power PC theory, Cheng (1997) has introduced a probabilistic measure of generative causal power, which is supposed to quantify the capacity of a cause to produce its effect. Of course, Cheng’s measure of causal power is not the only probabilistic measure of causal strength out there. Actually, there is quite a variety of different proposals in the literature (see, for example (Eells, 1991), (Suppes, 1970) or (Lewis, 1986)). But Fitelson and Hitchcock (2011) have convincingly argued that Cheng’s measure of causal power is the most suitable explication of intrinsic causal power, a concept that is highly valuable when it comes to prediction and decision making, since the intrinsic causal power of a cause is supposed to remain stable over different contexts. Additionally, Cheng and her colleagues have shown in several experiments that her measure is a very accurate description of how humans actually reason when it comes to causal relationships (see, for example, (Liljeholm and Cheng, 2007)). Despite all that, several arguments have recently emerged that challenge the adequacy of Cheng’s measure. Most notably, Sprenger (2018) argues that any measure that is not ordinal equivalent to Eell’s measure of causal strength (like Cheng’s measure of causal power) is deficient. By putting, what he considers to be, “a very general adequacy constraint” (Sprenger, 2018), called Generalized Difference Making (GDM), on measures of causal strength, he is able to show that only measures that are ordinal equivalent to Eell’s measure fulfill two very intuitive properties: Separability of Effects and Multiplicativity. While I agree with Sprenger that the violation of Separability of Effects and Multiplicativity amounts to a crucial flaw for any measure of causal strength, I do not agree that Cheng’s measure of causal power actually violates these two conditions. We have to consider that Cheng deduced the formula, which she employs as her measure of causal power, under assumptions that only hold in very simple causal scenarios. I will argue, that when we generalise Cheng’s power PC theory accordingly and make her measure of causal power applicable to more complex situations, it will actually satisfy Separability of Effects and Multiplicativity. Instead, Cheng’s measure does not satisfy GDM and I will argue that there is no good reason to think that this is a deficiency.

Bourgeois-Gironde, Sacha – see Salomone-Sehr, Jules

Božić, Aleksandar V.: “Explaining the vagueness of life: „individuals thinking“ vs. natural kinds approach”

In this paper I deal with the question of whether the vagueness of life can be better explained if life is conceived as an individual or as an instance of a natural kind. Life has been characterized as a phenomenon with vague boundaries (Malaterre 2010, Vlaardingerbroek 2012). This vagueness is both diachronic (concerning the historical transition from nonliving matter to first living entities) and

synchronic (concerning the borderline microscopic „grey area“ populated with entities such as viruses and macroscopic „grey area“ where the question emerges of whether ecosystems or the planet Earth are living entities in their own right). It has been proposed that the only familiar example of life, terrestrial life, is an individual belonging to a kind of „life-individuals“ with possible extraterrestrial instances (Hermida 2016) or that Life on Earth is an individual and not an instance of a kind (Mariscal and Doolittle 2018). This was criticized by Reydon (2019) who argues that both „individuals thinking“ and „kinds thinking“ can be appropriately utilized in a naturalistic metaphysics of biology. – My claim is that the vagueness of life is better explained with the notion of life as a natural kind. I will support this claim with the following argument: considering the diachronic vagueness, if prebiotic entities from which life gradually emerged constitute natural kinds then it is plausible that life also constitutes a natural kind. There is evidence suggesting that amino acids exist in the outer space (e.g., proteinogenic amino-acid glycine was found in the Murchison meteorite (Kvenvolden et al., 1970) and on the comet Wild 2 (Elsila, Glavin and Dworkin, 2009)). This entails a plausible assumption, i.e., that amino acids are a spatiotemporally unrestricted natural kind akin to chemical elements. Considering the synchronic vagueness, if entities at the border of life, such as viruses (with possible extraterrestrial instances, as presupposed by astrobiology), constitute a natural kind, then it is plausible that life also constitutes a natural kind. It is ontologically parsimonious to assume that the totality of living entities, emerging from kinds of prebiotic entities, possibly as different life forms on different worlds (or on the same world, as in the shadow biosphere hypothesis) form a natural kind. Vagueness of life can be explained with a diachronic and synchronic continuum between nonlife and life involving entities of various natural kinds, up to and including the natural kind of life. – The argument does not depend on an essentialist notion of natural kindhood. Moreover, I assume that a clustered notion of natural kindhood of life (such as the one proposed by Ferreira Ruiz and Umerez, 2018) can best account for both the diachronic and synchronic vagueness.

Brazil, Inti – see Malatesti, Luca

Buchholz, Oliver: “The Curve-Fitting Problem Revisited”

Deep neural networks (DNNs) are increasingly applied to tasks akin to the curve-fitting problem (CFP). However, while the fundamental statistical tradeoffs inherent to the CFP are well studied, it remains unclear whether existing results extend to the case of DNNs. In this talk, I argue for the opposite: DNNs escape the conventional analysis of the CFP. – In the first part, I outline the conventional analysis. The CFP refers to the task of fitting a mathematical function to given observations. This task is commonly taken to consist of two steps (Glymour 1981: 322): first, a general function class is fixed; second, specific values for the functional parameters are determined to choose the final function from that class. It is usually argued that *simplicity* as measured by the number of functional parameters should be achieved in the first step. In the second step, the goal is to maximize *accuracy* by achieving the closest fit to the observations (Turney 1990). According to the conventional analysis, there is a fundamental tradeoff between simplicity and accuracy (Forster and Sober 1994). Additionally, curve-fitting to compute predictions involves two different types of accuracy, namely *in-sample* and *predictive accuracy* (Forster 2002, Sober 2002). They differ in their relation to simplicity: high in-sample accuracy usually requires complex functions. Yet complex functions might be prone to fit idiosyncrasies of the given observations that are irrelevant for future observations and lead to poor predictive accuracy. This situation is known as *overfitting* (Hitchcock and Sober 2004). Simpler functions might prevent overfitting and achieve a higher predictive accuracy, however at the cost of lower in-sample accuracy. Thus in the CFP for prediction, the tradeoff between simplicity and accuracy becomes a trilateral relation in which simplicity mediates between two types of accuracy. – In the second part, I show how DNNs escape this

conventional analysis. From a curvefitting perspective, defining the overall architecture of a DNN corresponds to the step of fixing the function class. In a second step, an algorithm determines specific values for the network's weights, thereby choosing the final function from that class (Shalev-Shwartz and Ben-David 2016: 270). Just as in the CFP, the goal is to maximize in-sample accuracy in the second step and many DNNs even achieve an exact fit to the given observations (Zhang *et al.* 2017). Consequently, simplicity has to be sacrificed in the first step according to the conventional analysis. Indeed, DNNs are usually highly complex. Remarkably, however, research revealed that they exhibit high predictive accuracy regardless of their complexity and their very close fit to given observations. Put differently, DNNs are generally not susceptible to overfitting (Belkin *et al.* 2019, Poggio *et al.* 2020). Thus they seem unaffected by the tradeoff between simplicity and predictive accuracy, thereby escaping the conventional analysis of the CFP. – In the last part, I present a possible explanation for this peculiar behavior that was proposed in recent machine learning research. I conclude by discussing the result's implications for our understanding of the CFP and of statistics in general.

Buyse, Filip: “The Physiologist Johannes Peter Müller and the Philosopher Spinoza: An Underestimated Relation”

It is hard to believe that, in recent publications, nobody has systematically examined why “the father of contemporary physiology” quotes so explicitly from Spinoza's work, and refers to it at different stages of his impressive career. This is even doubly remarkable, given the fact that during the last decades there has been so much interest in Spinoza's philosophy among contemporary biologists (Antonio Damasio, Henri Atlan and Jean Pierre Changeux included) who argue convincingly that the Dutch philosopher (1632-1677) anticipated modern biological thinking. Likewise, it is amazing that Spinoza's name is completely absent in all the important biographies of Johannes Peter Müller (1801-1858). – This paper aims at filling in this striking gap by investigating the relation between Spinoza's sensory philosophy and Johannes Peter Müller's sensory physiology. After having analyzed the historical context with the role of Schelling and Hegel, it examines, in the second section, when and where precisely J.P. Müller mentions Spinoza (1632-1677) in his works. The idea that philosophy in general (and Spinoza's views in particular) did only play a significant role in an early stage of Müller's career will be contested. In a third section, it tries to find out - based on an examination of his early work *Über die phantastischen Gesichterscheinungen* (1826) as well as his magnum opus - *Handbuch der Physiologie des Menschen* (1837 & 1840) - why the professor at the University of Berlin applies the ideas of the Dutch philosopher rather than those of other influential early modern philosophers such as Descartes, Locke, Hume, or Hobbes. This part explores several elements of Spinoza's philosophy and claims that especially his views on the affections of the body and his innovative ideas on memory and hallucinations, were an important source of inspiration. Contrary to Piccolino & Drake (2013), this paper argues that in his revolutionary theory of sensations (with the law of specific nerve energies of the sense), Müller was directly influenced by Spinoza rather than indirectly from Galileo, whose ideas were transmitted via Kant and Locke. However, this paper defends the idea that also elements from Spinoza's double-aspect ontology were playing a significant role even though the 19th-century physiologist only seems to quote from his epistemology and his theory of emotions, being afraid to be accused of Spinozism, as he put it. Finally, it will be shown how Spinoza's views on animals having emotions inspired the comparative physiologist who quotes in this context from Spinoza's work. – Müller's main work *Handbuch der Physiologie des Menschen* (1837 & 1840) was in 1845 translated into French, and between 1838 and 1842 into English, so that his ideas spread out rapidly in Western Europe. Consequently, this paper will help not only to clarify the relationship between the influential Copley-medal winner and Spinoza, but also that between Müller and the myriad physiologists who were subsequently inspired by his work, Jacob Henle (1809-1885), Hermann Helmholtz (1821-1894), Rudolf Virchow (1821-1902), Theodor Schwann (1810-1882), Ernst Haeckel (1834-1919), Emil du Bois-Reymond (1818-1896) and Ernst

Wilhelm Ritter von Brücke (1819-1892), and their students such as Sigmund Freud (1856-1939), included.

C

Cevolani, Gustavo – see Peruzzi, Edoardo

Chikurel, Idit: “Maimon as a Baconian: Induction, Empirical Objects and Natural Histories”

In my talk, I assert that Salomon Maimon's (1753-1800) philosophy was very much affected by Bacon's work, and show that he was not solely influenced by the commonly mentioned philosophers such as Kant, Leibniz and Spinoza. Based on Maimon's commentary on Bacon's *Novum Organum* (*Bacon's von Verulam Neues Organon*, 1793), I discuss this influence in three main aspects: His use of induction, the employment of natural histories and his approach to empirical objects. Moreover, I show how turning to these three elements is intertwined with his skeptical stance towards necessary knowledge. For instance, Maimon employs induction to arrive at a higher degree of subjective necessity, a process that is infinite since, according to him, objective necessity of empirical knowledge cannot be achieved. Maimon also adopts Bacon's method of founding philosophy on the basis of natural histories. He presents a short history of mathematical inventions based on Montucla's *History of Mathematics* (1758), which includes many empirical discoveries, as well as a short essay describing philosophical systems, based on Bayle's *Historical and Critical Dictionary* (1697). Both histories serve as the grounds on which Maimon develops his own philosophical inquiry. This method embodies the idea of establishing knowledge on facts, not merely on symbolic cognition and ideas. In his regards to empirical objects, Maimon adopts a skeptical stance, since he believes that we may increase our knowledge of empirical objects and connections between phenomena, but we cannot show that judgments on empirical objects are objectively necessary. Accordingly, he rejects Kant's claim that judgments of perception can become judgments of experience and asserts that this transformation is impossible.

Crook, Barnaby: “The Compact Core - Emergent Structure Distinction in Artificial and Biological Neural Networks”

The renaissance of artificial intelligence, driven by the increasingly cheap availability of computation, is having a profound impact on scientific practice (Cichy & Kaiser, 2019). In the mind and brain sciences, deep artificial neural networks are increasingly being used as models of sensory and cognitive systems. For example, deep hierarchical convolutional neural networks demonstrate impressive performance on ethologically realistic tasks and unrivalled predictivity of the behavioural and neural responses of their target systems, such as the primate ventral visual stream and human auditory cortex (Kell et al., 2018; Yamins & DiCarlo, 2016). The success of this research program raises important questions. What kind of understanding do these models provide? And what kind of inferences are we licensed in drawing about their biological targets? – One influential theoretical framework, the deep learning framework for neuroscience, draws a principled distinction between the compact core of an artificial neural network, which includes its architecture, objective function, and learning rule, and the emergent structure that is learned by the network during training, instantiated in the weights and biases

between connected neurons (Richards et al., 2019). The scope of this distinction, which I call CED (Compact – Emergent Distinction), is the subject of my investigation. The methodological relevance of CED is clear; specifying the compact core is both tractable for scientists, and sufficient to produce highly predictive models, suggesting that more widespread adoption of the deep learning framework would be a useful way to make progress in modelling further sensory and cognitive systems. However, the philosophical implications of CED go further. Some authors suggest that the success of the deep learning framework justifies affording the compact core a privileged theoretical or epistemic status (Bashivan et al., 2019; Hasson et al., 2020). On this view, the success of the deep learning framework constitutes strong evidence that CED carves nature at its joints. For example, the CED could reflect the distinction between information encoded in an organism’s genome and information that the organism learns during its lifetime, or else the distinction between what is really teleological in a biological system, and what is not. – In this paper, I critically assess the scope of inferences that are warranted on the basis of the CED. Leveraging insights from the philosophy of biology, I argue that the CED does not carve biological brains at their joints. The sense of emergence which is exploited in the CED is an inherently perspectival property; one that depends upon which aspects of a system a research program is focused on (Callebaut, 2012). This means that taking different perspectives on biological systems will lead to different sets of properties being labelled as emergent structure. This may even include the objective functions, architecture, and learning rules that constitute the compact core in the deep learning framework for neuroscience. In order for our understanding of biological intelligence to benefit maximally from neural network models, we must take care not to overinterpret the scope of their properties.

Crüwell, Sophia: “Reframing the replication crisis as a crisis of inference”

The replication crisis describes a phenomenon in many empirical sciences, most famously psychology, in which several large-scale replication projects were and are unable to replicate the original results at all or only with much lower or reversed effect sizes. The cause for this crisis is largely seen in the intentional or unintentional misuse of statistical methods, combined with several cognitive and external biases. A further path to explaining or even fully accounting for field-wide replication failures is to consider the possibility of a low prior probability (or base rate) of true hypotheses in the field: if most or a substantial proportion of hypotheses that are tested in a field are in fact wrong, then a substantial proportion of positive results of significance testing will be false positives. The base rate of true positive results as a cause for replication failure has previously been considered by Bird (2020) and Ioannidis (2005). While considering the prior probability of a field’s hypotheses is important for understanding the replication crisis, the effects of biases, publication pressures, fraud, underpowered studies, overgeneralisations and a lack of formal theorising are undeniable. In this paper, I aim to give a better overall picture of the replication crisis by combining these explanations. To do so, I will give an extended Bayesian account of the replication crisis that centres the posterior probability of the hypothesis, i.e. the inference we make. I will take empirical evidence from the replication crisis, put this into the context of a Bayesian framework, and consider implications that follow from this. Specifically, I will argue that, in relevant areas of psychological research, the prior probability of the hypotheses is likely lower than generally thought, the likelihood of the evidence given the hypothesis is low due to small effect sizes and a lack of strong theory, and the marginal probability of the evidence is artificially large due to questionable research practices and biases. Our posterior belief in the hypotheses in many areas of psychology tested using standard research practices should therefore be weak when it is currently seemingly strong. Once we adapt our inferences accordingly, the large scale replication failures seen in e.g. social psychology will not be surprising anymore. I will conclude that, seen through an explicitly Bayesian framework, the replication crisis is better understood as a crisis of inference.

D

D'Alessandro, William: "Unrealistic Models in Mathematics"

Models are indispensable tools of scientific inquiry, and one of their main uses is to improve our understanding of the phenomena they represent. How do models accomplish this? And what does this tell us about the nature of understanding? While much recent work has aimed at answering these questions, philosophers' focus has been squarely on models in empirical science. I aim to show that pure mathematics also deserves a seat at the table. I begin by presenting two cases: Cramér's random model of the prime numbers and the dyadic model of the integers. These cases show that mathematicians, like empirical scientists, rely on simple (and often distorted or unrealistic) models to gain understanding of complex phenomena. There are also morals here for some much-discussed theses about scientific understanding. Two issues in particular are worth highlighting. First, modeling practices in mathematics seem to confirm that one can gain understanding without obtaining an explanation (contra [de Regt 2009], [Khalifa 2012], [Strevens 2013], [Trout 2007]). Second, these cases cast doubt on the idea that unrealistic models confer understanding by imparting counterfactual knowledge (contra [Bokulich 2011], [Grimm 2011], [Hindriks 2013], [Levy 2020], [Lipton 2009], [Rice 2016], [Saatsi forthcoming]).

Danese, Antonio: "Flowers and Teleology"

In 1862, Charles Darwin published a study, which seemed entirely dedicated to the interpretation of the morphology and biomechanics of the flowers of the orchid family, through a meticulous work of morphological analysis devoid of metaphors and the more theoretical aspects contained in *On the Origin of Species* (1859). – *On the Various Contrivances by which British and Foreign Orchids Are Fertilised by Insects, and on the Good Effects of Intercrossing* (Darwin, 1862; 1877) led to a close confrontation between different and opposing philosophies of nature, which implied, for those who adhered to them, a specific conception of the world that sometimes coexisted, but more often competed with, evolutionism in the attempt to interpret known and fresh plant phenomena. Hermann Müller, Fritz Müller and followers of the doctrine of secondary causes drew an armoury of facts from this contribution to vegetable physiology, wherewith to assail Paley and natural theology arguments, while Argyll, Federico Delpino, Asa Gray and all those who, entrenched in their impregnable bastion of faith and severely repelled chance in natural history, found new and marvellous instances of design in the pages of this treatise. – To clarify the connections among natural theology, design, teleology and Darwinian explanations about orchids, this paper departs from detailing the methodological differences between the author's studies of floral co-adaptations and traditional beliefs about their origin. I will focus on descriptions of *Coryanthes macrantha* and the study of homologies in *Catasetum* and *Malaxis paludosa* to show how Darwin initiated an inescapable action of eroding the consensus of botanists on the finalist model that informed Linnaeus' natural system. – Following this examination, I will refer to several contemporary interpretations of the meaning of Darwin's "flank movement" (Gray, letter to Darwin, 2–3 July 1862) to show the need to conceive contrivances according to a new conceptual framework where randomness, co-optation (exaptation: Gould, Vrba 1982), and rudiments (spandrels: Gould, Lewontin 1979) entail the replacement of teleological reasoning with a probabilistic approach in the study of nature where the hypotheses are tested through the study of homologies and selective dynamics. – The core of my argument is that Darwin's flowers allow for the disclosure of a methodological approach that is in radical opposition to explanations that resort to finalism. The study of the variety of forms and physiologies of orchids, in terms of coevolutionary selective processes that can emerge in the heart of random events, constitutes the fundamental Ariadne's thread that Darwin extends to the naturalists and philosophers of nature willing to follow him to the centre of the adaptive labyrinth of flowers.

Dardashti, Radin: "On the theory-ladenness of theorizing"

The theory-ladenness of observation or data is a much discussed topic in philosophy of science. It is common to consider the theory-ladenness as something problematic, which needs to be overcome to be able to confront theories with a neutral base. However, similarly theories themselves obviously are not being developed in a vacuum. So one might similarly ask the question whether there is a kind of theory-ladenness involved in theory development itself and whether that may pose a threat to the reliability of the theory, which is then not only motivated by the available empirical data. In this paper I discuss various kinds of theory-ladenness in theory development and the conditions under which they may or may not be problematic. – The theory-ladenness of theories can nicely be illustrated with examples from particle physics and theories of gravity. For instance, recent results at the LHC have disconfirmed many proposed theories beyond the standard model of particle physics. Many of these theories were motivated not solely for the purpose to account for yet unexplained empirical observations, but mainly to resolve certain theoretical problems. One very prominent example is the Higgs naturalness problem. The non-observation of these theories, however, has now put these theoretical considerations under pressure giving rise to more theory independent approaches like simplified models and the standard model effective field theory. Another example of this kind of less theory-dependent theorizing is Horndeski's (1974) theory, which besides General Relativity also includes Brans-Dicke theory and Quintessence as special cases. This has led to a more general confrontation of gravitational theories with recent results from experiments at the LHC and LIGO. – What purpose do these more or less theory-laden theories serve in the physical sciences? Should one always aim for more theory-unladen theories or is it epistemically advantageous to develop strongly theory-laden theories? I will argue that the answers to these questions can't be answered generically as the amount of theory-ladenness may serve different purposes in different circumstances. Nevertheless there are several general features that follow. In empirically inaccessible domains, theory-unladen approaches promise an epistemically more fruitful theory development (in agreement with Oriti (2019)). For meta-empirical approaches to theory evaluation, like the no alternatives argument by Dawid et al. (2015), it serves a precondition for the possibility of reliable meta-empirical theory evaluation. And finally more theory-unladen approaches allow to identify historically contingent features of theory development, which unnecessarily constrain future developments.

De Benedetto, Matteo – see Luchetti, Michele

de Bruin, Leon & Kostic, Daniel: “How evolutionary and environmental factors shape the relationship between structural functional connectivity”

The last decades have witnessed an increased effort in understanding the relationship between structural and functional connectivity models (respectively SC and FC hereafter) of the human brain. SC is modeled by using graphs in which nodes represent neural elements such as single neurons at the micro scale or neuronal populations at the meso and macro scales. Edges represent physical connections such as axonal projections between individual neurons at the micro scale or white matter tracts or fibers between neuronal populations at the meso and macro scales. FC is modeled by using graphs in which nodes represent blood-oxygen-level-dependent (BOLD) signals (in fMRI) or EEG channels (in EEG recordings) and the edges represent synchronization correlations between BOLD signals or EEG channels. The idea is that if two BOLD signals (or EEG channels) are synchronized, the populations of neurons that the BOLD signals represent (or brain areas that the EEG channels represent) are connected. – Early studies of SC and FC found that structurally connected areas of the brain exhibit a greater functional connectivity (Honey et al 2009). However, subsequent neuroimaging literature (Suarez et al 2020) shows that most functional connections are not supported by an underlying structural connection. In fact, functional connectivity is often higher in anatomically unconnected areas. Even when there are direct anatomical connections between functionally connected nodes, the correlation between structural

connectivity and functional connectivity is between $R = 0.3$ and $R = 0.7$ (Suarez et al 2020, 304). – The question is how we can make sense of the seeming mismatch between SC and FC. In our paper we argue that this can be explained by looking into the evolutionary and environmental factors that shape the development of structural and functional connectivity. We focus in particular on connectomic self-organization during the extended postnatal ontogeny of the human brain. As Changeux et al. (2020) point out, this connectomic self-organization can be understood as a Darwinian process of overproduction, stabilization, and elimination. In this process, at critical periods the connectivity configurations resulting from the growth cone wanderings produce a broad diversity of synaptic connections. Through synaptic pruning, this diversity is then reduced by the reciprocal exchanges of the developing brain with the outside world (Paquola et al 2019). This leads us to conclude that the evolution and development of FC are not determined by structure alone, but should be seen as a dynamical process guided by environmental interactions and adaptive pressure. We suggest that, from an evolutionary perspective, it makes sense to assume that the relationship between FC and SC is underdetermined to increase the organism's capacity to adapt and survive.

Desmond, Hugh & Ariew, André & Huneman, Philippe & Reydon, Thomas:
“The varieties of Darwinism: An integrated dynamic account” (Symposium
“The Legitimacy of Generalizing Darwinism”)

In the debate about how Darwinian ideas can be legitimately generalized, one core concern has been to search for an ideal-type Darwinian explanation which is generalizable to various nonbiological domains (Reydon and Scholz 2015; Schurz 2021). The idea is that, if such a generalizable template can be found, then this can ground judgments on what applications of Darwinism are legitimate and which are not legitimate. – A complication for this enterprise is how, in broader contexts, “Darwinism” does not only refer to an explanatory practice that helps achieve scientific-epistemological goals. Sometimes “Darwinism” is used as a near-synonym of naturalism. In other contexts, “Darwinism” is associated with “reductionistic” approaches to human behavior and history (e.g. Nagel 2012). In yet other contexts, is also used to demarcate scientific domains and scientific communities (e.g., sociobiology, evolutionary psychology, evolutionary economics, evolutionary history, etc.). In face of these broader usages of Darwinism, one can opt for two types of judgment. The first would be to lump the broader usages of Darwinism together into the category of “extra-scientific”: these meanings of Darwinism are thus added on, and are not essential to what Darwinism means. Darwinism is an ideal-type explanation (or a set of such explanations), and can be used for social purposes, both in the scientific community as well as in broader society. The weakness of this option is that it fails to make sense of why the type of explanation proposed by Darwin – and not, for instance, an equally universalizable theory, such as the second law of thermodynamics – was so amenable to generalization in a broader scientific and societal context. – In this contribution, we argue for an alternative view, where three distinct functions are essential for understanding Darwinism: an ideal explanation, a methodology, and a worldview. These functions influence each other. For instance, new applications of Darwinism-as-methodology can lead to revisions of the ideal type of Darwinistic explanation. We illustrate the interplay between Darwinism-as-worldview and Darwinism-as-explanation through historical examples. One is how Francis Galton saw that Darwin’s ideas “broke the spell of the old ‘argument from design’” (Radick 2019), and hence could be used for eugenicist purposes, i.e., to use principles of selection to “improve” the human species. The resulting account is a dynamic one, where scientific practices, scientific community structures, and political debates inform each other. In this way, Darwinism can be best thought of as a historical rather than as a natural kind.

Dolega, Krzysztof: “What can Free Energy Modelers Learn from Cybernetics?”
(Symposium “The Cybernetic Renaissance”)

The Free Energy Principle (FEP) is a formal framework that originated in computational neuroscience but has recently gained significant following across different fields of inquiry into mind and life. The eponymous principle at the core of the framework postulates that living organisms are a special subset of self-organising systems; or more precisely, that living systems ensure adaptive exchanges with the environment by maintaining their own states within a range prescribed by their phenotype. This is formally modelled as minimisation of the divergence between an organism's expected and actual states, which is equivalent to information-theoretic free energy. The popularity of the framework comes from its wide applicability and promise of delivering a unified account of perception, action, learning, attention, and planning (Friston 2010; Parr & Friston, 2017). – Many proponents of the FEP view it as a continuation of the cybernetic program (Pickering & Clark, 2014; Seth, 2015; Safron & Deyoung, 2020). Firstly, the two research frameworks share similarities in terms of their scope and focus on self-organising systems. Secondly, much like cybernetics, the FEP places special importance on feedback mechanisms, often stressing the circular-causality of action and perception. Finally, proponents of the FEP commonly appeal to concepts from cybernetics and control theory, such as Conant and Ashby's (1970) Good Regulator Theorem or Maturana and Varela's (1972) Autopoietic Theory, in order to explicate the main assumptions of their modelling framework. – However, despite multiple similarities, the FEP and cybernetics still differ from each other in some crucial ways. The first and most visible difference is the centralised nature of the free energy framework which, unlike cybernetics, is organised around a single formal principle. The second and less obvious distinction between the two research programs is in their adopted methodologies. While cybernetics embraced applied modelling, pursuing the construction of systems that implement some theoretical model to demonstrate its work in real-world settings (whether through Walter's tortoise robots [1950], the development of Rosenblatt's Perceptron [1957], or the Cybersyn project [Medina, 2006]), the FEP has predominantly focused on delivering theoretical models, i.e., models which demonstrate the feasibility of the formal strategies prescribed by its central principle (e.g., Perrinet, Adams & Friston, 2014; Constant et al., 2021). – The aim of this paper is to explore the different modelling strategies involved in the FEP and cybernetics, in order to show that many of the recent criticisms levelled against the FEP stem from the two differences mentioned above. Firstly, I will argue that FEP's singular focus on offering formal descriptions formulated in accordance with the central principle has hindered its ability to integrate already successful models of the phenomena that it purports to explain. Secondly, I will argue that in order to move forward, the FEP needs to deliver applied models of these phenomena. Although proponents of free energy modelling claim that the framework can deliver mechanistic explanations of many phenomena that lie within its scope, the most fruitful way to deliver on this claim is to embrace the real-world problem solving that initially motivated cybernetics.

du Crest, Agathe: “Objectivity at stake in mathematical models: the study case of evolutionary history” (Symposium “The Legitimacy of Generalizing Darwinism”)

At first, evolutionary theory was formulated in verbal terms by Darwin and Wallace. Only subsequently, with the Modern Synthesis and the especially work of Ronald Fisher, evolutionary theory was mathematically formalized. Many other mathematical formulations followed, as the Price Equation or Grafen's Formal Darwinism; and today weight is increasingly given to mathematics in other biological sciences, like behavioral ecology. – When it comes to transpose the Darwinian evolution to cultural and social phenomena, the question arises again about the role of mathematics in evolutionary explanation. What is the explanatory power of mathematical models? Can they be used for the purposes of generalizing Darwinism, where they are applied to any population likely to undergo evolution by natural selection? In order to assess their adequacy as « good » scientific explanation and, *a fortiori*, if they are better than verbal models, it is thus required to focus on the criteria held up in their favour. – In this presentation I critically examine two arguments Alex Mesoudi (in *Cultural Evolution*) uses to argue for

the superiority of formal (or mathematical) models over verbal models. The first is that conceptual (or verbal) models are faulted because of their *subjectivity*. This implies that a formal formulation is at best *objective* and at worst *less subjective*. I will rely on the case study developed by Mesoudi, from Peter Turchin's researches in history, about the successive cycles of rises and falls of empires in Europe, from 0 to 1900. I will propose the hypothesis that mathematical models are not able to achieve the assumption of objectivity underlying this argument. Moreover, that does not constitute a criterion on which their explanatory power should be based. It will lead us to wonder to what extent contradicting the argument affects other arguments in favor of mathematical formulations, especially Mesoudi's second one: their better ability to quantify reality.

Dustmamatov, Aznavur: "Geography as Science: The Limits of the Geo-Ontological Approach"

As a special science, geography faces the question of what is distinctive about its subject-matter, that is, whether geographic facts exist. There has been, in some recent scholarship, a tendency to equate the subject-matter of geography with entities found in its domain of study, e.g. lakes, tributaries, districts, borders, and the like. This tendency, which I refer to as 'the geo-ontological approach,' rests on two implicit assumptions: (1) if geographic entities exist, so do geographic facts; and (2) what makes a given entity geographic is its belonging to a distinct category of entities. As a result, the problem of the existence of geographic facts is transformed into the problem of whether geographic entities exist (Thomasson 2001). – I argue against both assumptions. First, the existence of a geographic object does not entail the existence of geographic facts. A lake has physical, chemical, biological, and other properties; whether any of its properties are geographic is a distinct question. Moreover, it might not be necessary for an object to exist, in order for there to be geographic facts about it; even fictional entities, or those that no longer exist, can possess what would appear to be geographic properties ('Sherlock Holmes's residence is in London'). – Second, what makes an entity geographic is not its belonging to a distinct category of entities. Following Tambassi (2017), I review various attempts to specify what this distinct category of entities might be: objects portrayable on a map (Smith and Mark 2001), mereotopological objects (Varzi 2007), mesoscopic objects (Smith 1998), the planet Earth or the planetary surface (Casati and Varzi 1999). Each proposal faces counter-examples that can be resolved only by appealing to the prior notion of the geographic in a circular fashion. I suggest instead that what makes some entities saliently geographic is their relationship to geographic facts, and, therefore, it is only by explicating the nature of a geographic fact that we can explain why some entities, and not others, are geographic. – I propose that there is a distinct level of facts that is geographic. The facts in this level emerge from the co-placement of distinct categories of entities (physical, chemical, biological, economic, etc.). They cannot be explained by appealing solely to facts intrinsic to the first-order entities, because any such explanation must involve the extrinsic relations that obtain among these components as a result of being co-placed. For example, how cacti survive with little water is a question that belongs to biology (plant physiology); however, why cacti occur in one area, but not in another, is a question of geography, for it cannot be answered without investigating what these biological entities are co-placed with, e.g. humidity and temperature. – In conclusion, I contend that my approach is preferable to subject-matter pluralism (Tambassi 2017), according to which there are multiple legitimate categories of geographic entities, for it preserves the unity of geography's subject-matter, while still acknowledging that there might not be a single distinct category of geographic entities.

E

Egg, Matthias: “Quantum Fundamentalism vs. Scientific Realism”

It is widely assumed that the only way to defend an ontologically serious brand of scientific realism about quantum mechanics (QM) is to go beyond the standard textbook account of QM and to opt for one of the versions (or interpretations) of QM that were developed in response to the measurement problem. This assumption seems to be confirmed by the observation that recent proposals which seek to remain neutral with respect to such interpretations (e.g., pragmatist, Bayesian or information-theoretic approaches) do not yield a substantive ontology of QM. – I will argue against this received wisdom by developing an account that takes standard (“textbook”) QM ontologically seriously, despite its failure to solve the measurement problem. I show that the customary association of quantum ontology with some particular solution to the measurement problem rests on a dubious identification of ontology with fundamental ontology and a concomitant failure to acknowledge the inherently non-fundamental character of QM. – After a general defense of my non-fundamentalist approach to ontology, I will demonstrate its viability in some concrete examples of QM. In particular, I will show how an ontologically robust view of some key concepts of textbook QM (spin, wave function collapse, wave packets) does not depend on choosing a particular response to the measurement problem. I will then reply to two possible objections: first, that the ontological posits of my proposal aren’t local beables, and second, that my account only purchases what looks like an ontological commitment at the price of semantic vagueness. –These considerations are of crucial importance for the wider question of how scientific realism can be applied to QM. Quantum theories have often been viewed as a stumbling block for realism, and although this may partly be due to a contingent historical association of QM with instrumentalism, there is also a more substantial reason for this view: since the choice between the different above-mentioned responses to the measurement problem is underdetermined by the empirical evidence, it is doubtful whether the scientific realist should commit herself to any one of them. – If, however, what I argue here is correct, then there is a substantial part of QM to which the realist can be ontologically committed, because it not affected by this kind of underdetermination. Admittedly, this is not a commitment to any fundamental ontology, but it still goes far beyond the purely empirical results of QM which also the non-realist would accept. I therefore conclude that a substantial scientific realism about QM is indeed possible, but only at the price of abandoning quantum fundamentalism.

Eytan, Yuval: “Hobbes on Scientific Happiness”

Because Hobbes was the first to offer ethical and political thought based on individual desires, many consider him the father of political individualism, claiming that his conception of happiness involved abandoning the ancient eudaimonic ideal and the Christian ideal of eternal happiness. In contrast, this article suggests that Hobbes’s conception of happiness derives from his understanding of scientific truth, thus adding an objective dimension to the accepted view that its foundation is strictly subjective and psychological. Highlighting previous commentators’ inattention to the link between truth and happiness in Hobbes’s thought, I demonstrate the inaccuracy of considering him the founder of a new ideal of happiness grounded in individual experience rather than an external knowledge, scientific or divine. – I believe Hobbes adopts the ancient principle that man’s happiness is necessarily conditional upon his submission to a normative system derived from the concept of truth. His originality lies in an innovative methodology for the study of truth. Bacon attributes to the new science a progressive element that can free man from a cyclic and eternal existence with no progress or significant change. I contend that the dynamic that Hobbes attributes to happiness must be understood in relation to the progressive nature he attributes to scientific knowledge. In addition, I assert that in this sense it is also essentially different from the dynamic he attributes to the preservation of life. The idea that every individual can be happy in this world expresses criticism toward Christianity, but originates in the optimistic notion that scientific conclusions can be reached in the realm of ethics and that rational and educated human beings will

choose to act in accordance with them. For Hobbes, of the many sciences intended to improve well-being, ethics is the most important, as it can provide a clear and precise definition of happiness, an objective definition that his predecessors failed to provide. – This interpretation sheds light on Hobbes’s understanding of progress as an objective element of his conception of human happiness. Everlasting happiness is made possible by subordinating inner faith to the timeless truth of the scriptures as mediated by the clergy. Worldly happiness may be realized by subordination to the scientific knowledge expressed in the laws of nature, which Hobbes hoped would be reflected in the laws of the state. In no way does Hobbes seek an ideal of life based on each individual determining by and for himself the purpose of his existence. For him, such a worldview is evidence of a highly negative state of affairs that humanity managed to escape by creating a new scientific method.

F

Fahrbach, Ludwig: “The abundance of scientific evidence for our best theories: Too much of a good thing?”

Scientific realism, the position that our best scientific theories like the theory of evolution and plate tectonics are probably close to the truth, is usually defended by appealing to the tremendous empirical success of these theories. Some realists have recently looked more closely at the actual empirical evidence that constitutes the tremendous empirical success. From this work a new difficulty for realism has emerged: The empirical evidence for our best theories is typically so plentiful that a proper assessment of the full details of the full body of evidence seems entirely infeasible for a layperson like a philosopher (Peter Vickers ms., *Identifying Future-Proof Science*, Alexander Bird 2017, “Scientific Realism and Epistemology”). Ironically, the very abundance of the evidence precludes a proper assessment of the evidence by a layperson. If this is right, a layperson has to rely almost entirely on trust in the testimony of the respective group of scientists. – In my talk I analyze the situation and aim to show that the problem is not as severe as Vickers and Bird contend. One kind of trust is inevitable, but unproblematic, namely trust that the reports of all the observations and experimental results relevant for a theory are largely correct. More interesting and relevant for the realism debate is the assessment of the inference from observation to theory. I argue that this task is not as hopeless as it seems. – First, the abundance of the evidence relevant for our best theories means that there is an enormous amount of *evidential over-kill*. Hence, a layperson does not have to examine the whole body of evidence, a suitable subset suffices. Call a subset of evidence *sufficient*, if one can infer the truth of the theory from it, while ignoring the evidence in the complement of the subset. When constructing a sufficient subset the layperson has a lot of choice which pieces of evidence she picks. She is allowed to “cherry-pick” the evidence (given certain conditions). She can pick pieces of evidence that are comparatively easy to understand, and ignore the most technical and complicated pieces of evidence. Furthermore she can pick pieces of evidence that offer especially strong support for or against the given theory. Or so I argue. – Second, most pieces of evidence relevant for a theory will involve a lot of detail. A layperson will only be able to understand mere sketches of the pieces of evidence. Vickers (ms.) worries that such sketches will be “full of holes” severely limiting the reliability of the layperson’s assessment. Against Vickers I argue that the limited reliability is counteracted by the *diversity* of the evidence. In general, a subset of evidence with much diversity can provide very strong support for a theory, even if each assessment of each kind of evidence in the subset has only limited reliability. This is intuitively plausible, but can also be backed by probabilistic arguments like the Condorcet jury theorem.

Feldbacher-Escamilla, Christian – see Baraghith, Karim

Feldbacher-Escamilla, Christian – see Gebharter, Alexander

Feldbacher-Escamilla, Christian & Haueis, Philipp: “Patchwork Approaches to Concepts and Different Scales”

In philosophy of science, patchwork approaches analyse how scientists use polysemous concepts with multiple related meanings (Wilson 2006, 2017, Novick 2018, Bursten 2018, Novick & Doolittle 2021). These approaches model polysemous concepts as patchworks with multiple patches, i.e. scale-dependent, technique-involving, domain-specific and property-targeting uses of a word. E.g.: in the domain of gases, “temperature” involves kinetic gas theory and refers to mean kinetic energy at the scale of molecules whereas in the domain of solids, “temperature” involves restricted ensemble approaches and refers to frozen order at the scale of polymer chains (Wilson 2017). In general, a patchwork concept is legitimate when its patches include techniques that are reliable, when its domains are homogeneous, and when each patch-specific property is significant to reach an epistemic goal. In this talk, we extend this general work on patchwork concepts by addressing hitherto unanswered questions about the notion of scale. – Patchwork approaches include scale to account for the “tyranny of scales”, i.e. the fact that many entities display different properties or behaviors at characteristic spatial, temporal or kinetic scales (Batterman 2013, Wilson 2017, Bursten 2018). However, this literature leaves important questions unanswered: (1) When does a change of scale generate a novel meaning? (2) Besides reliability, homogeneity and significance, what specific constraint governs concepts which have multiple scale-dependent uses? (3) And how can we relate multiple scale-dependent uses rigorously to one another? We answer (1) by claiming that a change in scale changes the meaning of a term if there are different discernible regularities about the behavior of the entities. Though not every change of scale in scientific inquiry changes the meaning of a concept, scientific concepts which change their meaning in a scale-dependent manner allow researchers to express more regularities about epistemically significant properties (construed as behaviors of entities but also as dispositions, mechanisms, or quantities). We answer (2) by introducing the matching constraint: the precision of a technique should match the scale at which an entity displays a property of epistemic significance. This constraint further clarifies the role of techniques in investigating scale-dependent properties and links scale changes to the epistemic goals associated with a patchwork concept. Measurement techniques need to be spatially precise enough to distinguish between two entities at the same scale, and temporally/energetically precise enough to capture regularities of the entities’ behavior that researchers aim to describe, classify or explain. To answer (3), we link the notion of scale in the patchwork literature to scales in the theory of measurement, such as the nominal, the ordinal, and the cardinal scale. This allows us to use measurement theoretical principles such as the construction of equivalence classes to bridge concepts among different measurement theoretical scales. Using “temperature” as example, our working hypothesis is that the quantitative (the temperature of x), the comparative (x is warmer than y), and the qualitative (x is warm) level can be construed as patches. Relating measurement theoretical scales thus may also be subject to the above-mentioned constraints of reliability, homogeneity, significance, and matching.

Fernández Pinto, Manuela – see Leuschner, Anna

Ferrari, Sacha: “Uberized science is the new black”

This early 21st century faces a severe skepticism toward science. Among these challengers, we can find various pseudo-scientific communities such as the Flat Earth Society, anti-vaxxers, astrologers, etc.

Besides these groups, some individuals decided to take part in the fight against orthodox science on their own. Their solitary practices include, for example, seeking online information about the reliability of Covid vaccines, medical auto-diagnosis by consulting a health forum (like Doctissimo.fr in the French-speaking community), building up a home-made experiment in one's garden to detect the curvature of the Earth, and creating a DIY biology lab in one's garage (Simons, 2021). This new way of inquiring information inside a social-media context can be seen as an uberization of scientific knowledge. Science is no more a matter of hierarchical verticality (with experts above lay people), but something horizontal and isonomic where information is produced, shared and sought from equal to equal (with individuals all interacting at the same level). According to this view, each of us is considered as an autonomous entrepreneur, a self-made and self-employed scientist who can run his or her own epistemological 'business' by himself in order to obtain reliable knowledge without relying on the blind scientific authority. This is the gist of the uberized science. – This talk aims to understand the causes of the emergence of this new epistemic strategy. We will argue that this uberization has been produced by (at least) three different factors: a technological one, a metaphysical one, and a political one. First, the rise of the Internet and new means of communication and information allowed a democratization and liberalization of the speech market (Aupers and de Wildt, 2021; Bronner, 2003). This opened a huge theatre stage where orthodox and newcomer heterodox scientific ideas are relentlessly struggling. Facing these epistemic battles behind a screen, each of us is free to choose their winner without fearing the peers' judgement of the outside world. Secondly, this new knowledge paradigm is the result of the disenchantment of our Western cultures induced by the failure of the 20th century Grand Narratives such as liberalism, communism, and Enlightenment (Lyotard, 1984). The positivist credo of the 19th century, promising progress of mankind by the help of science, is no longer credible. In light of its conflicts of interest (with Big Pharma for instance) and its legitimization of injustices (e.g. craniology), scientific authority appears nowadays illegitimate as the only reliable source of knowledge. Lastly, the rise of a new kind of neoliberal sociopolitical structure gave a political and metaphysical autonomy and independence to individuals. This new spirit of the capitalism started with new management techniques within companies in the nineteen-eighties (Boltanski and Chiapello, 2018) and expanded to the organization of the society itself. We will argue that the axiology of this DIY scientific inquiry relies on the same background as this neoliberal axiology (in its philosophical perspective at least). This paper demonstrates that uberized science is not an epiphenomenon but is a profound turnover of our society.

Fischer, Enno: "Naturalness: a Constitutive Principle"

Like the Kantian a priori a constitutive principle is a condition for the possibility of scientific theories. Unlike the Kantian a priori a constitutive principle may change in the course of scientific progress. Constitutive principles have gained prominence through Michael Friedman's (2001) work and, more recently, have been addressed from a pragmatist perspective with a particular emphasis on their role in scientific practice (Stump, 2015). – In this talk I aim to support and extend the pragmatist understanding of constitutive principles with a new case study: Higgs naturalness. The naturalness principle is widely believed to be an important guiding principle for research in high energy physics. One of its many formulations states that the physics at low energies should be largely independent of the physics at much higher energies. The Standard Model Higgs mass presents an apparent violation of this principle. If we assume that the Standard Model (SM) is valid up to the Planck scale, then the Higgs mass appears to depend sensitively on parameters that are located at a much higher energy scale. This is why many physicists expected that the SM is not valid up to the Planck scale and that new physics would be discovered in experiments currently performed at the Large Hadron Collider (LHC). But to date there are no conclusive signs of such new physics. – I argue that naturalness is best construed as a constitutive principle for current research practice in high energy physics. Current theories in high energy physics are construed as effective field theories. These theories provide an adequate representation of physics below a particular energy cutoff, independently of the physics at energies much higher than that cutoff. Naturalness is a condition for the possibility of such theories because it ensures the independence of the

energy regimes. Understanding naturalness as a constitutive principle has advantages over alternative construals such as naturalness as a guiding principle, epistemic value, or theoretical virtue. In particular, I will show that this helps us better explain why naturalness is considered to be so important by physicists (Giudice, 2019) and philosophers of physics (Williams, 2015), even if its status is put into question by the recent results at the LHC. – This result sheds new light on our understanding of constitutive principles more generally. First, constitutive principles may derive their main justification from an empirical basis and, second, scientific practitioners may have an explicitly pragmatic and fallibilist attitude towards the constitutive principles of their research practice.

Fischer, Stephan: „Zur Konzeption der Globalgeschichte“

Für die Wissenschaftstheoretikerin bedeutet es einen seltenen Glücksfall, Zeitzeugin einer konzeptionellen Neuerung einer wissenschaftlichen Disziplin oder eines Teiles davon werden zu können. In den letzten Jahren ereignet sich eine derartige Neu-Konzeptionierung innerhalb der Geschichtswissenschaft mit dem beachtenswerten Aufschwung der Globalgeschichte. Er ist angetrieben von „der Überzeugung vieler Historikerinnen und Historiker, [...] dass die klassischen Analyseinstrumente für eine adäquate Interpretation der Geschichte im Zeitalter der Globalisierung nicht mehr ausreichen“.[1] Sie möchte einem „methodischen Nationalismus und tiefsitzenden Eurozentrismus“[2] eine neue Perspektive entgegensetzen. – Der Aufschwung globalgeschichtlicher Forschungskonzeptionen wird von überraschend wenig Aufregung innerhalb der Zunft der Historiker begleitet, verglichen etwa mit dem Aufkommen der Sozialgeschichte. Das mag an der Zugkraft des Begriffes „global“ im Kontext der „Globalisierung“ liegen[3], hat aber einen bemerkenswert negativen Effekt: die mangelnde analytische Schärfe vieler der für die Konzeption zentralen Begriffe. Zu diesen gehören „modernities“[4], „transfer“[5], „exchange“, „(post)-colonial“[6] oder „nation“[7]. Sie bilden den zentralen Korpus der neuen globalgeschichtlichen Konzeption und sind doch weit weniger klar, als es scheinen mag.[8] – So sei etwa auf die Fragestellungen der „early-modernities“ und der „multiple modernities“[9] verwiesen. Eine genauere Analyse zeigt nicht nur die Unschärfe der Begriffe, sondern sogar Widersprüchlichkeiten bezüglich der Verwendung in den beiden oben genannten Konzeptionen. Dazu kommen normative Voraussetzungen, die der Intention, einem „tiefsitzenden Eurozentrismus“ zu begegnen zuwider laufen. Dahinter verbergen sich nicht zuletzt erhebliche Unschärfen in der Verwendung des – nur scheinbar eindeutigen – Begriffes „Europa“. – Auch die Begriffe des „Transfers“ oder des „Austausches“ ermangeln genauer Kriterien, die verdeutlichen könnten, ab wann Grade, Größenordnungen, Qualität oder Quantität von Transfer und Austausch tatsächlich globalgeschichtlich relevant werden. – Von den Resultaten der Analyse sollen hier insbesondere drei Aspekte vorgestellt werden. (1) stellt sich heraus, dass viele methodische Problemstellungen und Interpretationsfragen überhaupt nicht spezifisch globalgeschichtliche Relevanz besitzen. Vielmehr verweisen sie auf wissenschaftstheoretische Fragen, die an die Geschichtswissenschaft insgesamt zu stellen sind. Im hier untersuchten Zusammenhang steht dann aber die Frage im Raum, inwieweit solche Konzepte geeignet sind, den globalgeschichtlichen Ansatz tatsächlich methodisch als neu zu spezifizieren. – Eine solche Spezifikation wird (2) oft nur vorgeblich erreicht, wenn die Begriffe und Konzepte durch normative Voreinstellungen spezifisch globalgeschichtlich werden, Norm also vor methodischer Begründung steht. Dies führt einerseits zu der Frage wie genau die Fruchtbarkeit der globalgeschichtlichen Konzeption innerdisziplinär methodisch und an Erkenntnisgewinn orientiert motiviert und begründet, mithin der Anspruch einer neuen Perspektive methodisch fundiert werden kann. Andererseits wirft dies - an einer sehr aktuellen konzeptionellen Stelle - erneut die alte Frage nach Werturteilsfreiheit der Wissenschaft sehr konkret auf. – Schließlich werden (3) gewisse Konzepte mittels ihrer „heuristischen Fruchtbarkeit“ gerechtfertigt. Die genaue Analyse weist jedoch darauf hin, dass selbst für heuristische Zwecke die Schärfe der Begriffe deutlich zu unklar ist. – Insgesamt ergibt sich für globalgeschichtliche Forschungsprogramme ein erheblicher Bedarf an analytischer Schärfung zentraler Begriffe und Konzepte. Nur dann kann aus einer methodisch klaren Position und orientiert am wissenschaftlichen

Erkenntnisgewinn verdeutlicht werden, worin genau der fruchtbare Perspektivwechsel besteht und welche neuen geschichtlichen Erkenntnisse aus dem veränderten Blickwinkel resultieren können.

Fletcher, Samuel: “Causal Modeling as Counterfactual Semantics”

I propose to reduce interventionist definitions of causation to counterfactuals terms. The semantics for these counterfactuals use a variation on the Stalnaker-Lewis (S-L) semantics, adapted for application in science. (For simplicity, I focus just on deterministic causal models.) The S-L semantics employ the collection of possible worlds and a similarity ordering thereon (or some other equivalent structure), relativized to each world. Then, "if X were the case, Y would be the case" is true at a world w just when at all the worlds in which X is true that are most similar to w , Y is also true. – I keep the structural features of the semantics while replacing worlds with causal models. A causal model is a triple (V, S, A) , where V is a set of variables, S is a collection of structural equations for V , and A is a value assignment to the variables in V . Variables specify a range of values of a particular property, e.g., whether a switch is "on" or "off." A structural equation expresses the value of one variable in V in terms of a function of the others; e.g., $X=f(Y, Z)$ is a structural equation for the variable X in terms of variables Y and Z . There is at most one structural equation for any variable in V . The assignment A maps each variable in V to one of its possible values compatible with the structural equations. In contrast to some authors developing the causal modeling framework (e.g., Pearl 2009), I do not take the structural equations of S to represent unanalyzed, primitive causal mechanisms. They merely place constraints on the possible assignments A . This is important to enable an explicit definition of causal notions in terms of causal models. – Next I adopt a lexicographic similarity ordering on causal models that combines orderings on the values, structural equations, and variables. Similarity of variables takes precedence over similarity in structural equations, which takes precedence over similarity in values of variables. This allows one to evaluate counterfactual conditionals concerning (in both the antecedent and consequent) arbitrary syntactic combinations of variable assignments in the usual S-L way. – Next, one can represent an intervention through a modal operator – much like Pearl's (2009) do calculus – applying to particular variable assignments, one which is true in a causal model just when its variable assignment is true and the variable in question is exogenous. (A variable is exogenous in a causal model just when there is no structural equation for it.) Interventionist counterfactuals then involve different antecedents – e.g., "if do(X) were the case, Y would be the case" – than non-interventionist ones. To evaluate it at a model (V, S, A) , one must find the model most similar to (V, S, A) in which X holds and is exogenous. The above similarity ordering selects exactly the model arising from a surgical intervention on (V, S, A) , as that notion is usually defined in the causal modeling literature.

Fletcher, Samuel: “Replication is for meta-analysis” (Symposium “The Replication Crisis and Philosophy of Science” / replacement for cancelled talk)

The role or function of experimental and observational replication within empirical science has implications for how replication should be measured. Broadly, there seems to be consensus that replication's central goal is to confirm or vouchsafe the reliability of scientific findings. I argue that if this consensus is correct, then most of the measures of replication used in the scientific literature are actually poor indicators of this reliability or confirmation. Only meta-analytic measures of replication align functionally with the goals of replication. I conclude by addressing some objections to meta-analysis.

Forgione, Marco: “Feynman: Visualization and Understanding of Quantum Phenomena”

The paper aims at clarifying the type of scientific understanding of quantum phenomena that Feynman had when he was developing his famous diagrams. I will argue that Feynman’s understanding came from his capacity of having a partial visualization of quantum phenomena and that this visualization helped him writing the appropriate equations for the calculation of the transition amplitude. – Historically, the importance of partial visualization was already emphasized by Boltzmann (1974), for he believed that theories should provide pictures of the physical world that should guide scientific thought and experiment. These pictures, though, should not be considered as faithful representations (one-to-one) of physical phenomena. Similarly, Schrödinger believed that spacetime visualizability contributes to the making of good scientific theories – even though the theories do not represent reality. – How do these considerations apply to the case of Feynman diagrams? The answer is provided in De Regt (2017) where an extensive analysis of the relation between visualizability, intelligibility and scientific understanding is laid out. In such analysis, intelligibility is a “value that scientists attribute to the cluster of qualities that facilitate the use of the theory” (De Regt 2017, p.40). Visualizability is one of the qualities that contributes to making a theory intelligible; and a given phenomenon is understood scientifically if there is an explanation that is based on an intelligible theory. Feynman diagrams are then conceived of as a visualization tool that makes the theory more intelligible. – However, I maintain that Feynman did not develop the diagrams to make quantum electrodynamics more intelligible, but rather, that the visualization of quantum phenomena, i.e., the writing down of the diagram, was fundamental to the development of quantum electrodynamics. In support of this reading, we can consider some historical examples: (1) Feynman’s attempt to understand the Dirac equation is based on his previous works on path integrals, but it is also guided by the intention of having a physical system that would satisfy the equation: the quivering electron and paths counting. This example suggests that Feynman’s method of searching for a physical, partially visualizable system, traces back to his previous theories. (2) Feynman’s presentation at the Pocono conference was ill-received because it made use of diagrammatic simplification to avoid mathematical complexities and the audience was not used to such a visual-based thinking. The disappointing reception of his talk forced Feynman to publish his works in a mathematically more rigorous way. What emerges from these examples is that visualization was fundamental to Feynman’s scientific method and theory building and thus that visualization was fundamental to the development of the actual theory.

Frembs, Markus & Trotter, Frida: “Categorically classical: Lessons from no-go theorems in quantum foundations”

Quantum mechanics (QM) notoriously challenges some of our most deeply entrenched intuitions about the physical world: Does nature allow for ‘spooky action at a distance’? ‘Where’ is the moon when no-one looks? Is the cat dead and alive at the same time? Various no-go theorems in the foundations of QM suggest that nature is unavoidably quantum: they derive a conclusion from a parsimonious set of premises, which is then shown to be in contradiction with some aspect of the quantum mechanical formalism. Given the extraordinary extent to which the latter has been confirmed in experiment time and again, one is forced to give up at least one of these premises. As recently emphasised e.g. by Dardashti (2021), the different ways in which a no-go theorem is interpreted are directly correlated with the way in which one treats its formal components, one of which is the mathematical structure in which the premises are formalised. – Focussing on this aspect, the first part of this paper offers a unified perspective on a number of no-go theorems in quantum foundations by pointing to the inadequacy of classical concepts for modelling the natural world. While this perspective is not entirely new on an individual level (albeit expressed at times more, at times less clearly), here we stress its unificatory

power when considered across a variety of no-go-theorems—including the Kochen-Specker theorem (Kochen and Specker, 1967), Bell’s theorem (Bell, 1987), and the Groenewold-van Hove theorem (Groenewold, 1946). More precisely, by giving a precise definition of classical concepts in terms of category theory, we identify the formal aspects relevant for the derivation of each of these theorems. – Building on this result, in the second part we propose a broader philosophical program concerning QM as a physical theory: our proposal is to see QM as a conceptual toolbox, which provides us with the framework within which to think—and, accordingly, model—the natural world. In other words, QM plays a framing role for scientific thought by providing the fundamental conceptual tools to encompass in a unified picture our current empirical knowledge of the world. Among other things, this proposal suggests that QM does not directly postulate a fundamental ontology, as abundantly argued in authoritative monographs as e.g. Albert (1992), Lewis (2016), Esfeld and Deckert (2017), Ney (2021). Instead, the identification of a fundamental ontology becomes prerogative of more specific theories, as e.g. the standard model, built within the constraints of QM, but not directly of QM. – Finally, our argument opens the debate in quantum foundations to a wider community, both in physics, where a number of proposals have recently emerged attempting to reformulate QM in terms of category theory, and in philosophy, where our view of quantum theory may provide common ground for different discussions in the metaphysics and philosophy of physics, as well as in the philosophy of mathematics.

G

Gebharder, Alexander & Feldbacher-Escamilla, Christian: “Unification and Explanation: A causal perspective”

In this talk, we focus on two different views of unification and their connection to explanation from a causal perspective. Which account of unification gets things right and how unificatory power can be measured is still controversial. In this talk, we are especially interested two prominent approaches to unification:

Mutual information unification (MIU): A hypothesis has the more unificatory power with respect to pieces of evidence the more it renders these pieces of evidence (more) informative about each other.

Common origin unification (COU): A hypothesis unifies a body of evidence in so far as it posits a common origin for these pieces of evidence.

MIU has been defended by Myrvold (2003, 2017), and COU by Lange (2004). Myrvold stressed that according to a Bayesian decomposition, only "MIU contributes to incremental evidential support, and there is no scope, within Bayesian updating, for COU to add to the evidential support of the theory" (p. 93). And Lange claimed that "genuinely to unify [pieces of evidence], a theory must reveal them to have some deep common explanatory basis" (p.208) and that Myrvold's account is inadequate because it "sets the bar too low to distinguish genuine from bogus unification" (ibid.). – The main goal of this talk is to shed new light on both MIU and COU from a causal perspective. To this end, we draw on Reichenbach's insight that common causes screen off their effects (or render them less informative about each other in the presence of additional causal connections). Based on this simple idea, we propose a probabilistic measure for COU that in some sense complements Myrvold's probabilistic measure for MIU: According to this first probabilistic take on COU, a hypothesis has the more unificatory power the more it renders pieces of evidence uninformative about each other. As a next step, we will use causal Bayesian networks to represent different patterns of how a hypothesis can be causally connected to a body of evidence. As we will see, already focusing on the simplest causal patterns suffices to make some relevant observations. We apply Myrvold's measure for MIU and our first take on measuring COU to each of

these structures. The upshot of this will be that causal structure heavily constrains the performance of these probabilistic measures. – Next, we use the basic causal structures and our results about how the measures for MIU and COU behave to shed new light on the connection of unification and explanation in causal settings. It will turn out that both probabilistic measures of unification do a bad job as indicators for explanatory power. While the measure for MIU underperforms when applied to the elementary causal structures we discuss, the probabilistic measure for COU is too permissive. Based on this observation we further develop our probabilistic measure for COU by adding a causal constraint, which will improve its ability to indicate explanatory power significantly. We modify the measure for MIU in a similar way and compare it with the causal measure for COU.

Genin, Konstantin: “Causal Discovery and the Randomized Controlled Trial” (Symposium “Learning from Data: The Secret to Success”)

Randomized controlled trials (RCTs) are widely considered the “gold standard” of causal inference. Nevertheless, RCTs regularly come under criticism from Bayesians and theorists of experimental design. Recent justifications of randomization (Hernan, 2020; Deaton and Cartwright, 2018) focus on the fact that randomization ensures that average treatment effects (ATEs) are identified from experimental data and that standard techniques for estimating ATEs are statistically unbiased. But these arguments are too weak to justify RCTs. Other designs secure the same epistemic goods and, arguably, at lower ethical cost. If there is a (frequentist) justification for experimental randomization, it must lie elsewhere. – Suppose we are interested in whether a new therapy is an improvement over standard treatment. The treatment effect for a particular patient is the difference between the outcome that would occur under the new treatment and the outcome that would occur under the standard treatment. The average treatment effect (ATE) for the individuals in the study is the standard measure of the efficacy of the new treatment. Since treatment effects cannot be directly observed, ATEs cannot be straightforwardly computed. The theory of experiments is devoted to overcoming this difficulty. The ATE is identified if different degrees of efficacy of the new treatment give rise to different probability distributions over clinical outcomes. If the ATE is not identified, statistical inference is hopeless: an ineffective therapy would tend to give rise to the same kinds of observations as an effective one. Furthermore, a procedure for estimating the ATE is unbiased if, in expectation, the output of the procedure is equal to the ATE. One of the signal advantages of RCTs over observational trials is that a properly conducted RCT guarantees that the ATE is identified and that standard estimation procedures are unbiased. – Although randomization is sufficient to ensure that ATEs are identified and standard estimation techniques are unbiased, it is by no means necessary to secure these goods. If an experimenter is in a position to conduct an RCT, then they are also in a position to implement a design satisfying the instrumental variables assumptions given by Imbens and Angrist (1994). These are also sufficient for identifiability and unbiased estimation. Moreover, they are compatible with giving investigators and patients greater latitude in the choice of treatment. The standard arguments fail to single out RCTs as uniquely normative. – A successful frequentist argument for RCTs must focus not on identifiability or unbiasedness, but on efficiency. One estimator of the ATE is more efficient than another if, on average, it is closer to the ATE. A somewhat neglected literature (e.g. Wu, 1981; Li, 1983) focuses on demonstrating that randomized experimentation is the minimax efficient strategy. In other words: the worst-case efficiency of estimates of the ATE in the randomized design is better than the worst-case efficiency of estimates in other experimental regimes. However, these results depend on somewhat restrictive assumptions on the model connecting treatment and outcome. We consider whether generalizations could succeed in giving a uniquely normative frequentist argument for the purported scientific gold standard.

Gramelsberger, Gabriele – see Kasproicz, Dawid

Greif, Hajo: “Analogue Models and Universal Machines: The Separation and Realignment of Cybernetic Paradigms” (Symposium “The Cybernetic Renaissance”)

At the origins of the cognitive sciences, there were two distinct modelling paradigms that only seemingly and only partly map onto the contrasting research programmes of artificial intelligence and cybernetics: analogue and computer models. Whereas classical cybernetics and classical AI each largely followed one of the respective paradigms, connectionist AI and in particular the recent renaissance of cybernetic approaches in the cognitive sciences offer productive and strikingly diverse hybrid approaches to the construction and understanding of scientific models. The present contribution redraws some of the lines that connect the original modelling paradigms to these contemporary hybrid developments. – In his “Models and Archetypes”, Black (1962) introduces a concept of analogue models that offers a systematic but implicit account of the cybernetic models of his time, such as Phillips’ MONIAC (1950), Ashby’s homeostat (1952/1960) or Walter’s tortoises (1950; 1951). Analogue models are material models designed to offer perceptual or conceptual access to those elements and properties of a target system which are required for its explanation or understanding. The force of cybernetic analogue models in particular rests on the assumption of real isomorphisms between elements and relations in model and target system, in terms of both systems operating in accordance with the same set of ‘circular causal and feedback mechanisms’ (von Foerster et al. 1953). Designing a cybernetic model means to articulate a hypothesis concerning those mechanisms by means of a material structure. – Computer models, although forming a contrasting modelling paradigm, originated in the same broadly cybernetic context. Most notably, Turing’s (1936) work on computability provided the blueprint for digital computers and computer modelling. The key property of computer models is that they may establish any kind of modelling relation that lies within the mathematically delimited domain of computable functions – which implies a *prima facie* absence of the very constraints that make analogue models meaningful. Computer models are neither able to rely on analogies between concrete elements and relations in model and target system, nor on an assumption of shared mechanisms, nor on the ability of human observers to conceptually or perceptually grasp the pertinent analogies. The relevant elements, relations and properties will have to be determined in other ways, unless one assumes that the target systems are also fundamentally computational. – In sum, where analogue models recreate the specific structure and web of relationships of the original in some, a priori unspecified, different medium, computer models create or recreate a wide variety of structures and web of relationships in one *ab initio* specified medium. This is the methodological point at which the traditions of cybernetics and AI parted ways. By reference to the contemporary examples of Behaviour-based AI, Predictive Processing and the Free Energy Principle, this paper sketches the heterogeneous ways in which computational and cybernetic modelling approaches have come to realign. The distinction between analogue and computer models provides the historical background and the conceptual tool for carving out the unifying themes and the relatedness-in-difference in modelling approaches behind these diverse incarnations of the cybernetic renaissance.

Gressel, Céline: “The Usage of Extended Reality Technologies in the Contexts of a healthy Life and their Impact on Well-being”

With the increasing spread of Extended Reality (XR) technologies in recent years, new fields of application have also been increasingly opened up. One of these application fields is the broadly defined medical context. In the context of medical applications, there are various applications that use XR with the goal of achieving positive effects for a healthy and good life. This research around these applications deals extensively with technical discussions of feasibility, usability, and utility, but leaves out the question of the ontological distinctions between physical and virtual realities and their impact on

people's interactions and lives. In my talk, I will address these distinctions and show how features of XR and the resulting consequences for interactions in XR also affect physical reality and how they impact well-being and a good life. – After introducing the most important terms and principles of XR, I will show in which application areas which forms of XR are used in the medical context and how. In the style of a highly abbreviated state-of-the-art report, I will give the audience a quick insight into this large field. By presenting concrete technologies, projects and programs, I will not only show the manifold possibilities of using XR, but also explain the associated hopes on the part of medicine, which is mainly increasing health and wellbeing. In a second step, these hopes will be questioned from a sociological-philosophical perspective. Thus, the advantages of shifting interactions to XR over interactions involving physical presence will be contrasted with the disadvantages of the consequences of this type of interaction. I will show that all known approaches to using XR in medical contexts have the effect of dissolving, blurring, or shifting spaces and boundaries at very different levels. – If we now consider the impact of these developments on people's lives, the question is how the shifting of boundaries affect the perception of XR, other people and of well-being itself. Therefore, I reflect on the theory of subjective well-being as used in psychology. Following this theory, the fulfillment of desires is considered a central component of a good life. However, since there are desires that cannot be realized in physical reality, virtual worlds are seen as an important tool for desire fulfillment and therefore positively influencing well-being. Though, this raises the question of what impact it actually has on people's lives to shift needs that cannot be met in the physical world to XR, especially when the boundaries are becoming increasingly unclear. Drawing on Johnny Hartz Søraker's theory, "The Value of Virtual Worlds" (2010) I discuss the different aspects of well-being that must be regarded when talking about XR. Besides the question whether XR experiences are authentic and whether this diminishes their positive impact on well-being, I address the question of the right degree of wish fulfillment, the consideration of whether shifting wishes into XR leads to a shift of problems (namely away from working on the cause, towards alleviating the symptoms) but also questions about the consequences if wishes can be fulfilled in XR that would be forbidden in physical reality.

H

Halffman, Willem – see Kostic, Daniel

Harbecke, Jens: "Mechanistic Constitution as a Natural Law"

According to the "mechanistic approach", scientific explanation is generally not achieved by subsuming scientifically relevant phenomena under laws, but by analyzing the mechanisms that underlie a phenomenon or "constitute" it on several levels (cf. Bechtel and Richardson, 1993; Glennan, 1996; Machamer, Darden and Craver, 2000). On the basis of this basic assumption, intensive research has been carried out in recent years on topics such as the adequate definition of the constitutional relation (Craver, 2007b; Harbecke, 2010; Couch, 2011), the possible exclusivity of the mechanistic explanatory model and the possibility of other forms of explanation (Huneman, 2010; Chirumuuta, 2014) and the formal rules of constitutive inference (Harbecke, 2015; Gebharter, 2017a; Baumgartner and Gebharter, 2016; Baumgartner and Casini, 2017; Baumgartner, Craver, 2007a, 2007b; Krickel, 2018). – Against the background of this literature on the formal properties and methods of mechanistic constitution, in this paper I am concerned with the metaphysical status of mechanistic-constitutional relations. In particular, I want to understand whether at least some mechanistic constitutional relationships are laws in the metaphysical sense. In order to get an answer to these questions, I take as a starting point David Lewis'

general criteria for the lawlikeness. Lewis defined laws as those regularities that form axioms of the best deductive system ("BS"). The best system describes all phenomena there is, and it maximizes simplicity, informativity, and accuracy of fit. A known problem with this simple idea is that there are actually many ways to balance these criteria, and that it is impossible to determine whether systems can be compared effectively or accurately with one another without a transcendent measure. In light of this insight, modern BS theorists such as Loewer (1996, 2007), Cohen and Callender (2009), Albert (2012) and Frisch (2011) have tended to claim that simplicity criteria ultimately have a pragmatic dimension: The best system is the one that summarizes as much information about the world as possible in a way that is useful to us. – This pragmatic criterion also implies that the length of the evidence for deriving relevant theorems is an additional criterion for the BS. Based on this fact, Frisch and Cohen Callendar have concluded that supervenient non-fundamental laws have a good chance to find their way into the BS. – I agree with this general argument in my paper, but argue in addition to Frisch, Cohen and Callendar that at least some supervenient laws can be deleted from the BS if certain constitutional statements of law are included in the system. In addition, the explanatory benefit for us is often immensely increased by the inclusion of such laws, since constitutive relations offer reasons why certain superordinate generalizations or laws actually apply. – So my general answer to the question of whether mechanistic constitutive relationships are laws is, "It depends!" At least in some cases, the mechanistic constitution is lawful, and there are reasons to believe that some of these statements will enter the BS. However, by far not all such statements will receive this special status, despite being true.

Haueis, Philipp: "Patchwork concepts and the norms of explanation"

Patchwork approaches in philosophy of science aim to show that scientists use polysemous concepts with multiple related meanings to reach epistemic goals such as description, classification or explanation (Wilson 2006, Bursten 2018, Novick 2018). These approaches model polysemous concepts as patchworks with multiple patches, i.e. scale-dependent, technique-involving, domain-specific and property-targeting uses of a word. While general normative constraints on such patches have been articulated (Haueis 2021), it has not been discussed how using multiple patches of a patchwork concept in an explanatory text satisfies norms of explanation. In this talk, I focus on the concept of "force" in classical mechanics to show how patchwork concepts satisfy two explanatory norms: the ontic norm of completeness and the epistemic norm of intelligibility. – To see why "force" is a patchwork concept, consider how physicists apply it to billiard ball collisions: at the macroscale, "force" involves Newtonian kinematic models in the domain of discrete rigid bodies to refer to the coefficient of restitution. At the continuum scale, "force" involves partial differential equations in the domain of smooth continua to refer to internal stress waves. At yet lower scales, laminate or molecular modeling is used to refer to interfacial forces between laminae and action-at-a-distance between molecular bonds (Wilson m.s.). In multiscale models of collisions, "force" thus partially refers to more than one property (cf. Field 1973, Kitcher and Stanford 2000). Patchwork concepts with partial reference satisfy norms stressed by ontic and epistemic accounts of explanation. Ontic accounts hold that good explanations describe the (causal) structure of the phenomenon accurately and completely (Craver 2014), while epistemic accounts hold that they make the phenomenon intelligible to rational agents (Wright and Bechtel 2007). I first claim that using "force" in multiscale collision models fulfills the norm of completeness (Craver and Kaplan 2020) in a context-sensitive manner. For example: restitution coefficient and internal stress waves are sufficient if the balls do not crack internally, whereas interfacial forces and molecular forces become relevant when they do crack (Wilson ms.). Depending on physical context, different token explanations satisfy the norm of completeness by referring to different sets of scale-dependent properties underlying macroscopic collision phenomena. Second, I claim that patches of "force" referring to these properties also make the phenomenon intelligible to agents with limited computational capacities. Scale-specific submodels describe collision events using only few parameters, while avoiding other intractable physical processes (Wilson 2017). – The analysis of patchwork concepts in scale-bridging explanation

adds to the existing literature in three ways. First, it shows that accounts which integrate ontic and epistemic constraints on mechanistic explanation in the life sciences (Illari 2013) can be fruitfully extended to the physical sciences. Second, it shows that besides concepts which diachronically change their reference (Brigandt 2010), concepts which synchronically change their reference also contribute to explanatory goals scientists pursue. Third, the analysis reveals unnoticed convergences between multiscale modeling and mechanistic explanation: in both cases, explanations are fragmented, consisting of partially overlapping models with conflicting assumptions, rather constituting a unified and logically consistent theory (Hochstein 2016, Wilson 2017).

Haueis, Philipp – see Feldbacher-Escamilla

Hillerbrand, Rafaela – see van Panhuys, Marianne

Hirèche, Salim & Linnemann, Niels & Michels, Robert: "" (Symposium "Are all Laws of Physics Created Equal?")

Two approaches to elevating certain laws of nature over others have come to prominence recently: on the one hand, Lange (2007) and Yudell (2013) have argued that there are meta-laws which relate to laws as regular laws relate to particular facts, and on the other hand, Hendry and Rowbottom (2009), Tahko (2015), Hirèche et al. (2021a), and Hirèche et al. (2021b) have argued that some laws are necessary in a stricter sense than others. This paper is an attempt to clarify the relation between the two notions, as well as their applicability to physical laws. We will first argue that certain meta-law accounts (in particular that of Lange) are at the same time also non-absolutist law accounts. In a second part, we then argue that physical practice suggests — provided that laws are necessary in some sense in the first place — the need for a non-absolutist law account, but not for a metalaw account. (Among other things, we discuss paradigmatic examples from the notorious debate on the explanatory priority of symmetries vs. conservation laws.) Taking the results of both parts together, we arrive at the conclusion that only those meta-law accounts are viable which are also non-absolutist law accounts, and that they should be re-read as, first of all, non-absolutist accounts in any case.

**Hirèche, Salim & Linnemann, Niels & Michels, Robert & Vogt, Lisa:
"Scrutinising non-absolutist law accounts on physics: The case for a non-absolutist DTA account" (Symposium "Are all Laws of Physics Created Equal?")**

A common feature of all standard theories of the laws of nature is that they are "absolutist": They take laws to be either all metaphysically necessary or all contingent. Science, however, gives us reason to think that there are laws of both kinds, suggesting that standard theories should make way for 'non-absolutist' alternatives: theories which accommodate laws of both modal statuses. In this talk, we set out three explanatory challenges for any candidate non-absolutist theory and discuss the prospects of the two extant candidates in light of these challenges. We then develop our own non-absolutist theory, the essentialist DTA account, which combines the nomic-necessitation or DTA account with an essentialist approach to metaphysical modality in order to meet the three explanatory challenges. Finally, we argue that the distinction between kinematical and dynamical laws found in physical theories supports both non-absolutism in general and our proposed essentialist DTA view in particular.

Höhl, Anna Elisabeth: “Grasping and Explaining – The GE-Account of Scientific Understanding”

Scientists strive to understand the phenomena they are researching, and philosophers have recently picked up the task to analyze what scientific understanding is and how scientists achieve it. I present a novel answer to these questions, namely, the ‘Grasping and Explaining Account of Scientific Understanding’ (GE-account). – I argue that understanding is an ability and not a type of propositional knowledge. Understanding is an ability that is manifested by the two activities of grasping relations and articulating explanations. Grasping, in the sense of getting epistemic access to a phenomenon, enables scientists to realize that there is some relation the phenomenon stands in. Explaining this relation allows scientists to investigate the nature and details of this relation. In practice, understanding a phenomenon often is an iterative process consisting of several instances of grasping (aspects) of relations and articulating these in explanations. I illustrate this iterative process with an episode from biological research on zebrafish to understand the function of a specific gene for embryonic development. – The GE-account focusses on the understanding that individual scientists achieve of the empirical phenomena they are researching by employing scientific methods or practices and it accommodates the context-dependent nature of understanding. The specific historical and disciplinary context constrains which relations scientists can grasp and explain, because it provides them with specific resources. To understand a phenomenon, scientists must rely on the established background knowledge, accepted research skills (like experimental, mathematical, or modelling skills) and (material) equipment like investigative tools or computing capacity. As many abilities, understanding, too, can only be acquired by practicing it in a social setting. And to gain confirmation that one understood some phenomenon adequately, scientists present the results of their understanding, the resulting explanation, to their research community. By appealing to the community, the objectivity of the individual understanding is increased. – In short, I argue that the GE-account presents necessary and sufficient conditions for scientific understanding and it takes the following form: A scientist *S* has scientific understanding of an empirical phenomenon *P* in a context *C* if and only if

- i. *S* grasps (details of) relations that *P* stands in and articulates these relations in the form of new explanations of (aspects of) *P* (*manifestation condition*),
- ii. *S* possesses and uses (material) equipment, relevant knowledge and research skills provided by *C* and required for understanding *P* (*resource condition*), and
- iii. *S* is a member of a scientific community that enables *S* to understand *P* and that approves *S*’s understanding of *P* (*justification condition*).

The account specifies how the widely shared fundamental intuition that understanding is something like “seeing how things hang together” can be conceptualized for the scientific domain, an intuition that cannot be accommodated by any account that takes understanding to be a type of knowledge. While some knowledge is a prerequisite as well as a product of understanding, understanding itself is not a kind of knowledge. It is a cognitive ability to make sense of a phenomenon in a scientifically appropriate manner, an ability that can only be achieved by training and participating in scientific practice.

Hommen, David: “Poetry and Truth – Scientific Models as Perspicuous Representations”

Modeling is a central method in many scientific contexts. A key feature of models is that they provide idealized representations of their target system (TS), in the sense that they are comparatively abstract (certain properties of the TS are neglected) and/or distorted (certain properties of the TS are deliberately misrepresented). Given that science aims at true and complete theories about the world, however, the question arises how intentionally incomplete and inaccurate models of reality can be of any use in scientific theorizing. – According to cognitivist views (e.g., Contessa 2007), models are literally taken false, but, nevertheless, epistemically useful. They allow for “surrogate reasoning” (Swayer 1991),

i.e., by studying them, one can learn something about their targets. Depending on the specific view, surrogative reasoning is either enabled by a) an investigation of the similarities between the model and the TS, b) an investigation of the properties of the TS, or c) demonstrations performed in the model theory, the results of which are then interpreted in terms of the target theory. Either way, surrogative reasoning seems to presuppose the possibility of direct theorizing about a TS. Hence, on the cognitivist view, models turn out to be inferior to theories and ultimately futile. – According to noncognitivist views (e.g., Isaac 2013), by contrast, model representations are not only untrue, strictly speaking, but not even regulated by truth. They rather serve nonepistemic, i.e., pragmatic purposes – e.g., the generation of testable predictions, policy recommendations or ‘how possibly’ explanations. Yet, pragmatic success conditions seem to screen off modeling strategies from the metaphysical side of scientific theorizing. Models might be instrumental in generating useful predictions, recommendations, etc., irrespectively of whether they converge on true theories. – To steer clear of the dilemma between futility and instrumentalism, I suggest analyzing models as perspicuous representations in the sense of the later Wittgenstein (PI, § 122). A perspicuous representation offers a fertile new point of view of certain phenomena by describing them in analogical terms so as to exhibit certain features and sharpen our eye for certain connections. The understanding elicited by such analogies is, to a certain extent, poetic in the sense that those analogies cannot be replaced by others without altering the elicited understanding. It is conjectured that models in science function similarly as (quasi) poetic sense-making devices. – Models offer innovative and fruitful synopses of data that permit us to grasp what up until then may strike us as puzzling and unintelligible. Like a “liberating word” (PO, p. 165), models evoke meanings that elude and transcend the lexical meanings of the descriptions they employ. To the extent that their content is tied to their peculiar mode of presentation, that content cannot be paraphrased (let alone ‘de-idealized’) within other conceptual frameworks. On the proposed analysis, models are normative in that they ‘channel’ our experience and ‘guide’ our investigations. Nevertheless, they are constrained by how the world is and can be empirically put to the test. Hence, some models are objectively better than others, and our choice must be world-responsive.

Hörzer, Gregor: “Constitutive Relevance First: Mechanistic Explanations without Mechanisms?”

In the philosophical debate about scientific explanations in and beyond the life sciences, the approach of the new mechanists has taken center stage since the beginning of the 21st century (e.g. Machamer et al. 2000; Craver 2007; Glennan 2017). Based on the principle of descriptive adequacy, i.e. that an account of scientific explanations has to adequately capture the explanatory practices of scientists, new mechanists have set out to account for the fact that scientists frequently use the term ‘mechanism’ when they explain how the phenomena they are interested in come about. Although much progress has been made, and consensus in the literature about some aspects of how to best characterize mechanisms appears to be growing, there still exist a range of candidate characterizations that differ in the details, such as the complex-systems view and the acting-entities view (see e.g. Krickel 2018: Ch. 2), and on close inspection suggest a somewhat confusing background picture regarding the ontology of and relations between mechanisms, components of mechanisms, and mechanistic phenomena. Here, I argue that a careful examination of the explananda and explanantia and their interrelations within the new mechanist framework reveals that the notion of a mechanism is less central to the account of mechanistic explanations than it might prima facie seem. Some of the problems can be traced back to starting with the notion of a mechanism and then moving on to its components and the phenomena that mechanisms are supposed to account for. Focusing on some of the ideas recently discussed by Craver et al. (2021), I explore an account of mechanistic explanations that attempts to capture what the new mechanists have in mind but sets aside the question what exactly mechanisms are. Roughly, the line of thought is that

although new mechanists frequently claim that the explanantia in the mechanistic framework are mechanisms, on closer inspection it turns out that it is rather what is typically considered components of the mechanism that do the explanatory work by standing in constitutive relevance relations to the phenomenon that is to be accounted for: A composite system's activities are accounted for by activities of some of the composite system's components, and constitutive relevance relations single out those activities of the composite system's components that help to bring about the phenomenon in question (i.e., the working components) from those that do not play such role. Thus, the central notion in the new mechanist framework is not that of a mechanism, but that of constitutive relevance. – Of course, this does not mean that we should deny the existence of mechanisms altogether, or that characterizing the notion of a mechanism is of no use. However, focusing on constitutive relevance first, and properly distinguishing between epistemic and ontological accounts of it, without characterizing the notion of a mechanism beforehand can help to clarify some of the ontological problems that come with the more traditional new mechanist picture, and opens up new strategies in spelling out the details of the new mechanist framework.

Hoyningen-Huene, Paul: “Objectivity, the Ideal of Value-Free Science, and Rudner’s Objection”

The title of this paper that connects three topics leaves much leeway for its content. This is because there are not only various conceptions of objectivity defended by different authors, but also the claim that there are several equitable concepts of objectivity (for instance, Douglas 2004 and 2009; Reiss and Sprenger 2020). In this paper I shall defend the idea that there is an abstract and general core meaning of “objectivity”. What is seen by those authors as a variety of concepts or conceptions of objectivity is in fact a variety of criteria (or indicators) of or means to achieve one and the same objectivity, applied to a range of different things. The core meaning of “objective” derives from its opposition to “subjective”. Something is objective if it is free from distorting genetically subject-sided contributions. However, what counts as objective depends on the respective metaphysical context. For instance, colors are taken as entirely objective properties in one scientific context (for instance, what is the color of dinosaurs? Turner 2016), and as non-existent in others. I shall then discuss the ideal of a value-free science. Because this ideal has been disputed in the recent literature, I shall accept it only tentatively and discuss its relation to the objectivity of science. Given the previous discussion of objectivity, the status of the value-free ideal can be at best a criterion of or means for objectivity. This implies that the connection between the ideal of a value-free science and the objectivity of science is contingent. This in turn implies that any concrete violation of the ideal of a value-free science must be analyzed with respect to its consequences for the objectivity of science, because it is not a priori clear that such violations are detrimental to science's objectivity. Given this analysis, we can then turn to Rudner's famous objection against the value-free ideal because “the scientist qua scientist makes value judgments” (Rudner 1953). Contrary to Rudner, the problem does not only arise in the application of science, but already internally to science, especially in the decisions to perform potentially dangerous experiments. We can then discuss the question whether the value judgements that Rudner had in mind really pose a threat to the objectivity of science. The surprising result is that this is not the case. This is because social values that aim at preventing social damage due to a scientific activity or science's applications increase the level of the demanded objectivity of science. I shall conclude the paper with a discussion of yet rarely discussed detrimental influences of values on science. They concern moral restrictions on science and systematically biased choices of research problems. Here, social values indeed pose a serious threat to the objectivity of some scientific disciplines.

Huber, Lara: “Epistemic Significance: Broadening the Perspective”

When the Royal Statistical Society (RSS) finally launched its new magazine “Significance” in 2004, it became apparent that the name was chosen after much a heated debate, given that the latter, as a concept, even when associated with statistical analysis, very often is misapplied or misunderstood (RSS, 2004). Almost two decades later, we are not only confronted with a “replication crisis” in the experimental sciences but also, seemingly, a “credibility crisis” as regards the understanding and the evaluation of significance from an epistemological point of view. The question, what should be considered “significant” often is regarded as a matter of methodology in the first place, namely, if a given test statistics relies on a valid “significance level” (i.e., 0.05). – Statistical analysis differentiates—on methodological grounds—strictly between significant outcomes on the one hand and non-significant findings on the other. In recent years the threshold for defining statistical significance itself has been the object of a fierce debate. In response to growing concerns as regards the lack of reproducibility and under-powered studies, peers criticized existing standards of evidence for claiming new discoveries in research: for instance, in addressing the need to change the default P-value threshold for statistical significance from 0.05 to 0.005 for these kinds of claims (Benjamin et al. 2018). The latter, as is acknowledged, will neither fix nor compensate for malpractices such as reanalyzing data (“p-hacking”). – Taken together, threshold analysis impacts on the epistemic concept of “significance” by limiting its evaluation to issues of statistical methodology only. Drawing on key criteria of knowledge in science, the paper broadens the perspective on significance as an epistemic concept. Two accounts to significance are introduced—as regards the strength and scope of outcomes respectively: Whereas “evidential strength” refers to the degree to which a certain claim is immune to defeat, “evidential scope” addresses the import of findings that enable us to predict or assess the given reach of claims. The latter might focus on specific interests in research (e.g., modelling) or prospective uses of outcomes, for example in medicine (“evidence for use”; cf. Cartwright 2006). The paper reflects on necessary and sufficient conditions of significance as a valid epistemic concept—including its impact on the assessment of scientific claims in general (“significant truths”, cf. Kitcher 2001).

Huneman, Philippe – see Desmond, Hugh

I

Iranzo Ribera, Noelia: “Counterfactual Reasoning as Make-Believe”

The aim of this talk is to offer an account of counterfactuals in terms of *make-believe*, the type of imagination employed in Walton’s (1990) pretense theory of fiction, which has been more recently carried over to fiction views of models, most notably by Frigg (2010), Nguyen & Frigg (2020), and Frigg & Salis (2020). The motivation for the project is twofold: to accommodate the wide diversity of counterfactuals populating the counterfactual realm – they tell us about the relations between microscopic and macroscopic entities, the behaviour of idealised objects, etc. – and to improve on the shortcomings of traditional semantic analyses of counterfactuals. – Why counterfactuals? Counterfactuals are counterfactual conditionals with *nomologically impossible* antecedents; they have a looser connection to the features of the actual world than other counterfactuals which are only contingently so. Despite their appeal to nomological impossibilities, they pervade scientific practice: they inform scientific explanations, reasoning about superseded theories, and above all else, model-based reasoning (Tan 2019, pp. 37-38). – Finally, they aggravate the prospects for semantic analyses of counterfactuals proposed by Stalnaker (1968) and Lewis (1973), which have traditionally been used to assign truth-values to counterfactuals. I argue that there are several problems with these accounts. In a nutshell, they

offer a bad reconstruction of counternomic reasoning in practice, possible worlds contain too much information for successful counterfactual evaluation (Salis & Frigg 2020), and the implicit epistemology of modality that results from these semantic accounts is underdeveloped (Kment 2006, Roca-Royes 2012). Furthermore, and specific to counternomics, when conjoined with the view that laws of nature are metaphysically necessary possible worlds semantics output vacuous counternomic truths, as there are no possible worlds where the laws are violated. – For the above reasons I propose to turn to the underexplored strategy of examining counternomics through the lens of *fiction*. The notion of fiction here employed is *fiction as imagination*, where imagination is understood as *pretense* in terms of games of make-believe à la Walton (1990). Briefly, games of pretense are initiated by props, objects or texts which *trigger* some direct imaginings, which together with principles of generation (PGs) prescribe additional imaginings. Applied to counternomics, make-believe results in the following framework: antecedents become prescriptions to engage in *legitimate* pretenses, where legitimacy is dictated by the functional role these counternomics play in scientific inquiry. The unificatory power of fiction is thus highlighted by the account's focus on this capacity of counternomic antecedents to prescribe scientifically legitimate imaginings. – Regarding truth, counternomics are fictionally true iff their consequents are prescribed for imagination by their antecedents and pertinent PGs. In the case of counternomics, these PGs are mostly logical rules of inference, and actual principles and laws of nature other than the one/s being suspended or altered in the antecedents. Hence, I defend that the functional account of fiction sketched is not incompatible with metaphysical underpinnings: despite the fact that in these imaginary scenarios some laws are altered or suspended, counternomics are partially made true by features of the actual world.

J

Jukola, Saana: “Bodies of Evidence – Determining the Cause of Death and the Problem of Underdetermination”

This paper addresses the problem of underdetermination as it relates to one of the central tasks of forensic medicine, namely determining the cause of suspicious deaths. I show how determining the cause of death involves numerous potentially value-laden judgments, which are partially influenced by the institutional and cultural context where the pathologist operates. I present the case of Excited Delirium Syndrome (ExDS) as an illustration of why dealing with this challenge is critical. – I begin the talk by outlining how underdetermination understood as an epistemological problem concerning the relation between data and hypotheses (Longino 1990) is salient in the practice of forensic medicine. As the debate on the COVID-19 death rate demonstrates, determining the cause of a death can become controversial matter even in circumstances where no foul play is suspected (e.g., Pappas 2018; Amoretti & Lalumera 2021). When a suicide, homicide or terminal occupational illness is a possibility, causal inferences related to past events and the death become even more contested. It has been argued that this is partially due to practices in forensic medicine not being evidence-based or standardized (Meilia et al. 2018). Findings from autopsy and biopsies have to be interpreted in the light of existing pathophysiological knowledge, toxicology, epidemiological studies, genetics, applied physics etc. (Meilia et al. 2018). Moreover, contextual factors, such as the witnesses' interpretation and reports of the mental state of the deceased individual, influence the interpretations of the physical evidence. This often leads to disagreements between experts. – In the second part of the talk, I will analyse the debate about ExDS, a controversial diagnosis often used in cases of deaths on police custody (Strömmer et al. 2020). In particular, I examine the role of background assumptions (Longino 1990) in influencing the production and interpretation of evidence that has been used for arguing for and against the existence of the condition since Charles Wetli and David Fishbain first suggested the diagnosis (Wetli & Fishbain 1985). In this case, different assumptions related to pathophysiological knowledge, circumstances of

deaths and relevance of experimental knowledge enabled drawing conflicting conclusions from the same physical evidence.

Jurjako, Marko – see Malatesti, Luca

K

Kant, Deborah: “Deep peer disagreement in set theory”

Against the backdrop of an unresolved debate on deep peer disagreement, I study a case of disagreement from the set-theoretic research context. Set theorists can prove for many statements that they are independent of their standard theory, ZFC. Adopting new axioms would prove some of these statements, but not all set theorists favour this direction. Based on the results of an interview study with set theorists, I model a situation, in which a pluralist and an absolutist disagree on the following framework propositions:

- P1: Independence results are final answers.
- P2: ZFC suffices and there is no need for further axioms.
- P3: No new axioms will be adopted by the community.

The pluralist believes P1, P2, and P3, while the absolutist believes their negations. I argue that this is a case of deep disagreement according to [Fogelin(2005 (1985))]. Further, I suggest that, considering the data, it is plausible to assume that the pluralist endorses the Pluralist epistemic principle: If 'pm' is proven independent from ZFC, then you are justified in believing that 'pm' is not about a matter of fact. And the absolutist endorses principles like the following. Absolutist epistemic principles:

- If 'pm' solves many questions outside of set theory in the way that practitioners predict and leads to a coherent theory, then you are justified in raising your degree of belief that pm.
- If 'pm' implies generic absoluteness of some important theory, then you are justified in raising your degree of belief that pm.

Then, the disagreement is also deep according to [Lynch(2010)]. As a third aspect, the pluralist and absolutist are modelled as epistemic peers according to a notion of peerhood in the scientific context that is built on [Cruz and Smedt(2013)]. – Having established these premises, I address the question of the rational response in such a situation. [Lougheed(2018)] claims that remaining steadfast is rational, while [Matheson(2018)] defends the equal weight view, according to which the pluralist and the absolutist should suspend judgement. Since the pluralist and the absolutist qua set theorists both aim at set-theoretic progress, rationality is first evaluated according to this group epistemic goal. It is shown that their framework beliefs influence their research, but that this does not lead to separate sub communities, i.e., they agree for substantial parts on what counts as set-theoretic progress. I argue that remaining steadfast on the framework beliefs is epistemically beneficial, because diverse mathematics is defended to be more fruitful than restricted mathematics. On the other hand, suspending judgement is probably epistemically detrimental. The arguments mainly rest on aspects of the set-theoretic community. However, if we evaluate the situation in terms of the individual epistemic goal of true beliefs about the framework propositions, the situation changes. From this case study, I conclude that the distinction between different epistemic goals as well their hierarchy matters to the debate on deep peer

disagreement, and that the apparent disagreement between [Lougheed(2018)] and [Matheson(2018)] actually is none.

Kaspowicz, Dawid & Wenz, Daniel & Gramelsberger, Gabriele: “How to Explore Scientific Code? (in Philosophy of Science)”

From a mere helpful auxiliary tool the computer developed into an integral part of scientific research. It is not merely a helpful device but forms part of what science is and does. How does algorithmic thinking change science and technology? How do scientific concepts transform in this new environment? To tackle such questions, we need new tools that help us to explore scientific code and identify structures that make it possible to explore new concepts and methods. In our presentation, we give an overview on this new field and introduce a software tool to tackle these problems. – The transformation mentioned above concerns not only the empirical part of scientific research but also the theoretical core of scientific reasoning. Our working hypothesis is that the translation of classical mathematical formulations as well as the translation of scientific concepts into the more restricted form of programming languages changes significantly not the only the objects but the subject of science itself. The idea is that today you need to grasp the workings of the code in the same way as you need to be able to find your way around the mathematical apparatus employed by elementary/particle physics to say anything useful about the nature of electrons. We also suspect that, as new models and new forms of mathematical representations can give rise to new concepts and objects in scientific theory, the same goes for new and different code and programming languages. – We will start with an analysis of the situation and state our general task: How to locate well known and how to extract new and relevant concepts from scientific code to make them approachable for the philosophy of science? We will also clarify how our tool-based method differs from mathematically oriented approaches like (Lenhard 2017) or epistemological like (Winsberg 2010). We then will list the problems that our approach faces, from inner theoretic obstacles like the transformation of mathematical into numeric representations (Gramelsberger 2011) to problems like code-readability. We will identify the latter as one of the most urgent problems: Most of the code scientific researchers produce is so called spaghetti-code - code that is written just to work, without any considerations about traceability from a third party. In many instances, this results in code that is only traceable by the research team itself. Finally, we introduce our software tool to improve the access to and the readability of the code for scholars from philosophy of science. GICAT gives not merely visualizations of the code structure that helps with code-readability, but it is foremost an analytic tool for concept formation and contextualization. It uncovers functional dependencies, keeps track of inheritances, depicts the dependencies on external libraries and works as a general comment extractors.

Kästner, Lena: “Multiplexes: New Directions for Computational Psychiatry?”

What does it take to diagnose, explain, treat and prevent mental illnesses? A common answer to this question is that we need to understand their nature and causes. But what does that amount to, precisely? Clinicians, scientists and philosophers have been seeking to develop models and theories of mental illnesses for centuries. While some research traditions have focused on the phenomenological aspects of mental illnesses (e.g. de Haan 2020), others have been looking into neurobiological substrates (e.g. Shelton 2007, Goodkind et al. 2015) and genetic underpinnings (e.g. Wong et al. 2008, Avramopoulos 2018). – Nowadays, the view that mental disorders are best understood as brain disorders is quite prominent (e.g. Insel & Cuthbert 2015, Walter 2013, Kandel 2018). It fits naturally within the medical tradition of seeking a common (molecular) cause for various symptoms to diagnose an illness as well as with popular naturalist-reductionist views of the mental. However, the brain disorder view is increasingly coming under pressure; for its exclusive focus on a single organic substrate seems too narrow (Adam 2013, Kendler 2009). Indeed, understanding mental illnesses requires looking at a variety of different factors contributing to the development and persistence of, as well as the recovery from,

mental illness. That is to say, scientists must take into account the role of, e.g., behavioral, psychological, neurophysiological, genetic, pharmacological and environmental influences on psychopathology. – Driven by the rise of computational methods on the one hand and the availability of large amounts of real-world data in psychiatry on the other, a whole range of mental disorder models have recently been suggested to come to rescue. Among them are multiplexes (e.g. de Domenico 2017, Braun et al. 2018, de Boer et al. 2021). Multiplexes are essentially networks of networks. As such, they can integrate data from different factors, at different scales, or across time. Intuitively, these multi-layered networks structures present quite appealing models of mental disorders that can be constructed by powerful computational machinery based on increasing amounts of real-world data. – In this paper, I systematically assess the potentials and challenges of multiplexes by comparing and contrasting them with other (simpler) psychopathology models. My examination will highlight that multiplexes actually face a range of challenges familiar from other species of psychopathology models. In order to get off the ground, they require answers to important foundational questions such as what variables and relations to include in a model and how exactly multiple different factors can be linked. Pragmatic and heuristic assumptions about systematic variable relations may help constrain multiplexes and thereby address this issue—but these assumptions are not inherent to multiplexes. Still, if properly constrained, 4D multiplexes may provide new directions for psychopathology research.

Katic, Ana: “The Dynamical Biological Explanation: A New Perspective for the Concept of Superorganism”

The traditional explanation of the concept of superorganism in biology is based on the notions of individual selection and Hamilton's rule. Considering only these two biological notions, this explanation fails to fully explain the key characteristics of superorganism. For instance, in the case of eusocial insects, Hamilton's view based on the haploid-diploid genetic system fails to explain how the degree of the relative relationship between sisters in a colony is greater than between mother and daughter. Also, using the new technologies it was found that some types of insects do not possess such a (haploid-diploid) genetic system, yet they are eusocial species. A new model that would completely explain the organization of the colony, based exclusively on the concepts of Hamilton's rule of individual selection and kin selection – has never been offered. This gene-based explanation is interpreted as a reductionist approach to the phenomenon in the philosophy of biology. The main assumptions of a reductionist approach are a belief in a) a linear hierarchy of the causal powers, that is a gene predominance (bottom-up causal relation), and b) that a natural selection operates only as an individual selection. We find the problem with this kind of explanatory strategy is in its theoretical assumptions. These assumptions are one-sided and cannot provide a detailed and fruitful analysis of the concept of superorganism. It is therefore inevitable that, following this gene-based explanation, we conclude that the concept of superorganism has a narrow and limited use in biology, presenting only a heuristic agent and the essentially replaceable hypothesis. We offer a new, dynamical type of explanation of superorganism, constructed by using the systems theory in a particular biological context. This kind of explanation is based on the idea of an organism as dynamic self-organization. The main assumptions of our approach are that a) the biological units are functional, thermodynamically open wholes, and complex systems with non-linear causal powers (no basic elements, such as genes, exists with all causal power, but causal relations are interconnected and mutually dependent/top-down causal relation), and b) a natural selection operates at every level of selection, like an individual selection as well as a group selection. In other words, epigenetic factors and the dynamic in a specific environment – factors which reductionist considered negligible – we find essentially important. The advantage of our explanatory approach is that it reveals the insights into microevolutionary processes, such as the emergent communication within the ant colony (the problem that seemed to be an enigma in biology for a long time), as well as into macroevolutionary processes, such as Gaia scenarios (fundamentally dynamical and non-linear self-

sustainability of the biosphere where the biodiversity is represented as an inherently autocatalytic process on every level - as a subset or as a whole). We conclude that the concept of superorganism has a multidimensional, robust and potent, both theoretical and experimental, significance in biology and astrobiology.

Khosrowi, Donal: “Extrapolating Causal Effects - Where Is Our Theory of Confidence?”

Extrapolating causal effects is a widespread epistemic activity in biomedical and social sciences. It involves measuring the causal effect of an intervention in some study population and endeavours to predict the effects of the same or a similar intervention in a distinct target. Extrapolation is difficult, however, as study and target settings often differ in important respects. Correctly predicting an effect in a target hence requires a host of empirical assumptions pertaining to relevant similarities and differences between settings and these assumptions, in turn, need to be empirically supported. – While the existing literature has made important progress on outlining strategies for extrapolation, the role of uncertainty has received rather little attention: causal knowledge is often scant and underdeveloped and crucial assumptions will routinely remain in need of additional support. Two questions, in particular, are in need of attention: first, how can we express our uncertainty and confidence concerning specific causal assumptions? Second, how do uncertainty and confidence compound and propagate onto a causal conclusion? – To make progress on the first question, I consider a bayesian networks approach (Bovens and Hartmann 2003; Landes et al. 2018; Poellinger 2020). While bayesian networks help us compute probabilities for specific causal hypotheses given diverse evidence, they cannot tell us how different assumptions work together in yielding a conclusion. The relationships between causal assumptions are not (merely) evidential: we need to consider which assumptions are necessary to make specific inferences; how relevant they are individually; and how they interact, e.g. whether they are causally, logically, or probabilistically related. – To make progress on the second question, we must hence consider how relevant our assumptions are for a conclusion and how they work together in enabling it. To help with this, I sketch a hybrid causal-graph based approach, called support graphs. Support graphs involve three layers. The first encodes causal knowledge and assumptions: structural causal models and corresponding graphs encode our amalgamated knowledge about the phenomena of interest and the assumptions we need for an inference. The second layer is a support layer. Drawing on a bayesian networks approach, it encodes how available evidence bears on the assumptions contained in the first layer. The third layer is a relevance layer. It encodes how relevant causal assumptions are for a specific conclusion, which can be investigated by performing sensitivity analyses, i.e. comparative static changes to the causal model/graph at the first layer to learn how a conclusion changes under these manipulations. Together with information from the second layer, performing such analyses also allows us to map out which conclusions enjoy how much confidence (in the spirit of Roussos et al. 2021). – In sum, support-graphs can help extrapolators clarify several interconnected issues: 1) which causal assumptions are needed for an inference, 2) how relevant these assumptions are for a conclusion, 3) whether they enjoy sufficient support, and 4) how confident we may be in certain kinds of conclusions. In virtue of these promises, a support-graph approach can facilitate our ability to articulate, manage, and ameliorate uncertainties in extrapolation.

Koenig, Daniel: “Objectivity and Subjectivity of Mathematics. On the Status of Mathematical Objects in Ernst Cassirer's Philosophy of Culture.”

In the history of philosophical reflection, there has been much argument about the status of the objects of mathematics: How can necessary truths be formulated about mathematical objects while at the same time being applicable to a reality determined by contingency? – In the research on the work of Ernst

Cassirer, who is primarily known for his early epistemological and late cultural philosophical writings, mathematics predominantly plays a role only insofar as it serves as a paradigmatic example of his turn from Substanzbegriffen to Funktionsbegriffen. However, Cassirer's extensive writings about the philosophy of mathematics in a narrower sense, which accompanied him in each of his creative phases, has received less attention and has been less illuminated. In the sense of Kant, Cassirer also deals with the question of the status of mathematical objects already in the context of his early epistemological works, insofar as he also ascribes objective, necessary validity to mathematical propositions. In my lecture I would like to point out that the interrelation of truth and act of thinking is crucial: Cassirer emphasizes in "Substanzbegriff und Funktionsbegriff" (1910, ECW 6. [SuF], translations are my own) that the totality of true propositions is always related to pure, non-empirical acts of cognition, just as these always aim at objective truths. Objectivity is thus not founded in a reference to something existing in itself, but in the accomplishment of the act of thinking itself: Mathematical truths correlatively correspond to an "activity of thinking, [which] is strictly regulated and bound activity". With this correlation of subjectivity and objectivity, the necessary and objective validity of mathematical propositions is tied to the "functional activity of thinking, [which] finds its hold in an ideal structure of thought" (SuF, 341). – Finally, in the context of Cassirer's turn toward a philosophy of culture, the question of the relation between objectivity and subjectivity undergoes a decisive pluralization, insofar as Cassirer arguments first for a dependence of physical facts on principles in the spirit of Heinrich Hertz. From here on he speaks of other forms of objectification that are based on principles other than those constitutive of the forms of scientific cognition. According to the basic thesis of the third volume of the "Philosophie der symbolischen Formen" (Erster Teil: Die Sprache (1923), ECW 11, Zweiter Teil: Der Mythos (1925), ECW 12, Dritter Teil: Phänomenologie der Erkenntnis (1929), ECW 13), his main work in the philosophy of culture, a deeper understanding not only of theoretical cognition in general, but also, according to the thesis pursued in the further course of my lecture, of the objectivity of mathematics is to emerge from the overall view of all symbolic forms, in their mutual difference and functional unity. By means of an analysis of the specifically mathematical forming of symbols or signs it will be shown that the propositions of mathematics continue to assert their status of necessary, objective validity, but are at the same time bound to a certain form of thinking, i.e. to a certain form of "subjectivity".

Kohár, Matej: "The Scaling-up Problem from a Mechanistic Point of View"

The scaling-up problem (more often referred to as the representation-hunger problem), has been a mainstay in philosophical treatments of non-representationalist accounts of cognition. Representation-hungry cognition includes phenomena in which the behaviour of a cognitive system seems to be sensitive to absent or abstract properties of the environment. In this paper, I will argue that employing the mechanistic framework has the potential to resolve the most pernicious aspects of the scaling-up problem and bring the non-representationalist on level ground with the representationalist with respect to scaling-up. This is because the mechanistic framework allows the non-representationalist to make use of mechanistic compositionality analogously to how representationalists make use of semantic compositionality when scaling-up representational theories of content. – The paper proceeds as follows: I first introduce the scaling-up problem in its familiar guise as representation-hunger. I then consider currently available responses to representation-hunger and argue that they are mostly based on a misreading of the representation-hunger problem as a conceptual issue about certain cognitive phenomena. I argue that such conceptual reading is unwarranted, and that the representation-hunger problem should be viewed as an open challenge to provide non-representational explanations of representation-hungry phenomena. I further examine some such attempts known from the literature and argue that even providing explanations of individual representation-hungry phenomena is not sufficient to completely dispel representation-hunger. This is because the non-representationalists still lack a heuristic for generating novel possible explanations for any given representation-hungry phenomenon. – Further, I argue that the scaling-up problem arises not just for non-representationalists, but also for

representationalists. However, representationalists have an available solution in semantic compositionality. Basic representations can compose to create representations of absent or abstract properties. Researchers can utilise this compositionality to generate possible explanations of arbitrary phenomena which involve sensitivity to the abstract and the absent. Popular non-representationalist theories of cognition, on the other hand, do not have any compositional explanatory posits that could be used in a similar fashion. – Then I introduce the mechanistic framework and show that mechanisms exhibit a form of compositionality. Mechanistic compositionality is the result of the hierarchical organisation of mechanisms and the fact that the properties of higher-level phenomena depend on the components and organisation of the constitutive mechanism. I argue that this mechanistic compositionality allows us to resolve the scaling-up problem for non-representationalists, because it also allows one to create how-possibly models of representation-hungry phenomena. This is possible, because we can predict how tweaking, replacing or reorganising some components of the mechanism would affect the phenomenon constituted by that mechanism. Together with the better-known strategy of functional decomposition, mechanistic compositionality provides the required heuristic for generating possible explanations of arbitrary representation-hungry phenomena. – From a mechanistic point of view, then, non-representationalism has the same resources for resolving the scaling-up problem as representationalism. Both approaches can rely on the same compositional strategy.

Kostic, Daniel – see de Bruin, Leon

Kostic, Daniel & Halfman, Willem: “Explanatory imperialism: empirical evidence for the claims about pervasiveness of ‘mechanisms’ in the life sciences”

The literature on scientific explanation in the philosophy of science has been dominated by the idea of mechanisms (Craver and Darden 2013; Bechtel and Richardson 2010; Glennan 2017). The basic idea can best be captured by the following definition of a minimal mechanism (Glennan 2017, 17): A mechanism for a phenomenon consists of entities (or parts) whose activities and interactions are organized so as to be responsible for the phenomenon. The new mechanist philosophers often claim that all explanations in life sciences are mechanistic in the above sense, or at the very least that they conform to various degrees of completeness of this definition, e.g. there could be full-fledged mechanisms, partial mechanisms, or mechanistic sketches (Boone and Piccinini 2016; Piccinini and Craver 2011). Furthermore, anything that doesn't fit this definition, or a degree of completeness thereof, is not an explanation at all (e.g. Craver 2016). We call this set of claims “explanatory imperialism”. But such extraordinary claims require extraordinary evidence, which so far wasn't forthcoming. The importance of empirical evidence about pervasiveness and uses of “mechanisms” in life sciences is particularly needed because examples and case studies that are used to illustrate new mechanists' claims cannot represent a statistically relevant sample, even if taken all together. Furthermore, given that they are admittedly handpicked, a robust quantitative and qualitative scientometric study of the large body of relevant literature that we conduct in this paper will put such claims into perspective by showing: (1) To what extent exactly do uses of the notion “mechanism” and “mechanistic explanation” conform to the accepted definitions of mechanisms and mechanistic explanation in the philosophy of science literature? (2) What is the pragmatics of uses of these notions that do not conform to the accepted definitions of mechanisms and mechanistic explanation in the philosophy of science literature? – In conducting this study, we will employ the following methodology. In the first step, we will create a corpus of high-impact articles in the field and algorithmically search through them. The second step will be the qualitative analysis of different uses of the notion of mechanism detected in the corpus. The purpose of the qualitative analysis is to classify these uses. Finally, we argue that the proposed methodology will provide comprehensive and empirically grounded insights into the debate on explanatory imperialism.

Koutroufinis, Spyridon: “The Phenomenon of Organism – Three Different Levels of Analysis”

The term ‘organism’ defines, together with the term ‘life’, the biological field of study. One can distinguish between two inherently different philosophical methods of approaching the concept of organism: Under the term phenomenal analysis of the organism, I subsume approaches to the nature of the organism that are limited to the description of the phenomena essential for living beings. In contrast, what I call the causal-ontological analysis of the organism targets the theoretical justification or explanation of the phenomenal analysis. – This lecture focuses on the ‘phenomenal analysis of the organism,’ from which I distinguish three levels. In each of the successive levels, the topic is treated within a higher level of abstraction: (1) Organism as the body of a living being: In the first level of phenomenal analysis, the term ‘organism’ denotes the body of a living being. This implies an understanding of what a living being is, i.e., the essential differences between living and inanimate beings, which in turn requires an understanding of what ‘life’ means. Since antiquity, this problem has been met with ever-increasing lists of the essential characteristics of living entities. However, as phenomenologically oriented philosophers have claimed, this approach presupposes an intuitive, non-discursive approach to life that is anchored in our empathy with our own being alive. Hence this level of phenomenal analysis is the less abstract one. – (2) Metabolism as the most essential feature of organisms: A change of focus from the living being’s material constitution to its self-continuation through its own processuality distances the concept of organism from the continuously changing living body and elevates it towards the principle that governs a living being’s continuous material transformation. Thus conceived, the term ‘organism’ becomes a principle of identity that refers to a nexus of reciprocally conditioning processes, which maintains or varies its own material form through the continuous exchange of the nexus’ material constituents. By focusing on the principle governing the exchange of matter, this definition elevates metabolism to the most essential feature of the organism. – (3) Emphasis on the autonomy of organismic self-organization: There is a rapidly growing body of publications in the fields of genetics, epigenetics and developmental plasticity that demonstrate the ability of cells and multicellular organisms to autonomously manipulate and reorganize their morphology, physiology and genomes down to the molecular level of the latter. This reorganization enables the organism to radically restructure their bodies so that, even under extreme internal and external circumstances, they can create most of the conditions for the continuation and, if necessary, the targeted change of their metabolism. This ability is the most basic biological faculty that is ubiquitous in the realm of life. The depth of organismic self-organization cannot even approximately be attained by the most complex examples of physical or chemical self-organization in inanimate matter. The autonomous manipulation of the conditions that determine the internal functions and external actions of organisms identifies the uniqueness of the organismic mode of being.

Krickel, Beate: “What mechanisms can do for (the philosophy of) cognitive science and psychology other than explaining” (Symposium “Mechanisms in the Cognitive and Social Sciences”)

Contemporary philosophers of science agree that mechanistic explanation plays a crucial role for cognitive science and psychology (Wright and Bechtel 2007; Piccinini and Craver 2011). Psychologists as well as cognitive scientists, among other things, are interested in explaining how humans think, reason, solve problems, make decisions, imagine, categorize, perceive, and so on. Depending on the particular research project, these how-questions are answered by describing the mechanisms underlying the cognitive capacities at different levels of abstraction. Cognitive neuroscientists usually describe mechanisms in of neurobiological terms, while cognitive scientists as well as psychologists usually

describe mechanisms in functional terms. A theoretical framework that highlights the relevance of mechanistic explanation for the special sciences has become popular under the label new mechanistic approach. One important insight of the new mechanists is that the mechanisms referred to in mechanistic explanations are individuated in terms of the phenomena they explain (Illari and Williamson 2012). In practice this means that the investigation of mechanisms must inform how phenomena are individuated (Craver 2007 Chapter 4.4). For example, it took a while for scientists to realize that the folk-psychological category “memory” indeed corresponds to different underlying mechanisms. Consequently, it is now assumed that memory corresponds to a number of different phenomena, such as working memory, episodic memory, and semantic memory. – In this talk, I will extend this latter insight. I will argue that the mechanistic framework cannot only be used to make sense of how cognitive capacities are explained. The resources of the framework—correctly understood—are much richer: based on a recent proposal for an epistemic account of mechanistic explanation, I will show that we can derive a scientifically grounded ontology of cognitive capacities. Roughly, I will argue that cognitive capacities should be individuated based on their categorical bases. The categorical basis of a capacity, I will argue, is best understood as the minimally necessary components of all mechanisms that underly manifestations of that capacity (i.e., behaviors). As an outlook, I will show how this mechanistic ontology of cognitive capacities provides a starting point for a fresh look on old questions such as (i) Is cognition/the mind extended? and (ii) Can there be unconscious mental states?

L

Leuschner, Anna & Fernández Pinto, Manuela: “Research on Shooting Bias: Social and Epistemic Problems”

As many philosophers of science have argued, diversity within scientific communities is epistemically valuable as it leads to a broader range of criticism. However, there are limitations to this epistemic benefit. In particular, the denial of well-established scientific findings, such as anthropogenic climate change or gender bias, can come with social and epistemic costs when it meets epistemic and political asymmetries in society and contributes to a social atmosphere that is hostile to science (Biddle/Leuschner 2015; Leuschner/Fernández Pinto 2021). Thus, it seems justified in some cases to exclude certain voices from the exchange of opinions. – In this paper, we’ll explore this problem further in light of a new case study: research on shooting bias. The shooting bias hypothesis aims to explain the disproportionate number of minorities killed by police. We’ll first present the mounting evidence that supports the existence of a shooting bias among police officers, especially but not exclusively in the US. Then we’ll focus on two studies by James et al. (2016) and Fryer (2016) who have claimed that—although they corroborate widespread racism in non-lethal police use of force—they cannot confirm a shooting bias. While we grant the authors good intentions, we consider the studies highly problematic: The authors have made questionable generalizations and presented the studies in a way that made it easy for right-wing groups and media to misuse them. Not surprisingly, the studies have been embraced and disseminated by powerful right-wing media outlets, such as Breitbart and Fox News, and white-supremacist websites, such as stormfront.org. Consequently, the studies have been both epistemically and socially detrimental as they have contributed to a social atmosphere in which anti-racist campaigns and scientists working on relevant topics have been attacked. However, in contrast to studies that bluntly deny well-confirmed scientific findings, such as the existence of anthropogenic climate change or gender bias, the situation seems more complex here. We’ll argue that the shooting bias studies could have been socially and epistemically useful if the findings concerning the existence of shooting bias were more carefully interpreted and communicated. We’ll undergird our argument by drawing upon Kitcher’s “Millian Argument against the Freedom of Research” (Kitcher 2001, ch. 8) as well as recent debates on epistemically detrimental dissent.

Lin, Qiu: “Du Châtelet on Mechanical Explanation versus. Physical Explanation”

In her second edition of the *Foundations of Physics*, Du Châtelet advocates a three-fold distinction of explanation: the metaphysical, the mechanical, and the physical. While her use of metaphysical explanation (i.e., explaining via the Principle of Sufficient Reason) has received some attention in the literature, little has been written about the distinction she draws between mechanical and physical explanations, including their demand, scope, and use in physical theorizing. This paper aims to fill this void, arguing that making this distinction is a crucial piece of Du Châtelet’s scientific method. According to Du Châtelet, a mechanical explanation is one that ‘explains a phenomenon by the shape, size, situation, and so on, of parts’, whereas a physical explanation is one that ‘uses physical qualities to explain (such as elasticity) ... without searching whether the mechanical cause of these qualities is known or not’. I will analyze Du Châtelet’s views regarding (1) What counts as a good physical explanation, (2) Why a mechanical explanation is not necessary for answering most research questions in physics, and (3) Why a good physical explanation, instead, is sufficient for answering those questions. In so doing, I argue that Du Châtelet is proposing an independent criterion of what counts as a good explanation in physics: on the one hand, it frees physicists from the methodological constraint imposed by mechanical philosophy, which was still an influential school of thought at her time; on the other, it replaces this constraint with the requirements of attention to empirical evidence, for that alone determines which physical qualities are apt to serve as good explanans.

Linnemann, Niels – see Hirèche, Salim

Livanios, Vassilis: “Thin Powers and the Governing Problem”

A metaphysically loaded version of the Inference Problem for laws requires a metaphysical explanation of how a second-order external nomic relation “descends” to the first-order level of property-instantiations producing a regularity. Recently (2021), Ioannidis, Livanios and Psillos (ILP) have argued that any adequate solution to the Inference Problem should explain how nomic relations manage to determine the “behaviour” of their properties-relata. The associated problem is what ILP dub the Governing Problem. – ILP discern various possible grades of modal strength of fundamental properties, from those properties that have the power to confer on their bearers a fine-grained disposition to “obey” the exact form of a relevant law (thick powers) to those that have the highly generic power to be nomically governable in general (ultra-light powers). In between these extremes, there exist other grades of modal strength of properties which ILP collectively call light powers. – ILP suggest that a necessary move to address the Governing Problem is to embrace the view that properties have light or ultra-light power, that is, what they collectively call thin powers. They argue however that the appeal to thin powers is not sufficient to solve the Governing Problem. Having a thin power can at best be conceived only as a partial ontological ground in virtue of which properties could confer dispositions on the objects that bear them. As ILP argue, the other indispensable parts of the full grounds of dispositions are the nomic relations among properties. Like the thin powers, nomic relations cannot by their own solve the Governing problem because they, qua external, can relate their specific relata (and consequently govern their “behaviour”) only if the relata have thin powers to be nomically relatable. – Given that, ILP have introduced a dualist model, according to which nomic relations and properties with thin powers are individually necessary but only jointly sufficient in order to have an adequate metaphysical explanation of the actual behaviour and dispositions of objects. The dualist model however needs further elaboration

because it leaves some crucial questions unanswered. In this talk, I focus on one of those questions: do actual fundamental properties have light or ultra-light powers? I argue that, due to the light-powers-view's assumption that properties can confer certain "substantive" (but nonetheless thin) dispositions to their bearers, the view in question faces at least two significant problems. The first problem is that, generally, there is no specific "substantive" disposition that can be associated with each property and each law in which that property appears. Compared to the ultra-light-powers-view which posits a unique disposition for all thin powers, the light-powers-view is then ontologically less economical. The second problem is that any choice of a thin but "substantive" disposition that the advocates of the light-powers-view might make would be arbitrary. Given these two problems, I conclude that the defender of the dualist model has reasons to prefer the ultra-light-powers-view over the light-powers one.

Lopez, Luis: "Machine Learning Models and Understanding of Phenomena"

The deployment of Machine Learning (ML) models in scientific research has shown that they can make accurate predictions in domains where traditional models or simulations have failed to do so (e.g., AlphaFold 2 and the protein-folding problem). However, science is not just about prediction (or pattern recognition); it is also about understanding (de Regt, 2017). In this talk, I address the following questions: can Machine Learning (ML) models provide understanding of phenomena? If so, how? And, more importantly, what is the nature and reach of that understanding? I argue that the answers to these questions depend on whether we are talking about interpretable ML models or opaque ML models. Here, I follow the distinction made and defended by Rudin et al. (2021). Namely, while an interpretable ML model "obeys a domain-specific set of constraints to allow it to be more easily understood by humans," an opaque (or black box) ML model is a "formula that is either too complicated for any human to understand, or proprietary" (ibid.). (I am not concerned with proprietary black boxes in this contribution). I show that this distinction has significant implications not only for understanding the inner workings of the model itself (as it directly follows from the short definitions given above) but for the understanding of its target phenomenon. Thus, contrary to Sullivan (2019), I argue that model opacity –in the context of ML– impairs understanding of phenomena. Moreover, I show why link uncertainty reduction, achieved through empirical evidence supporting the link between the model and the target phenomenon, cannot compensate for the negative effects of opacity. I also demonstrate why the opacity of non-interpretable ML models cannot be treated as implementation black boxes. To illustrate my point, I focus on Deep Neural Networks (DNNs) –which not only represent the quintessential black box but are also the focus of Sullivan's paper– and compare them with traditional scientific models used in model-based explanations. In addition to the opaque DNNs examined by Sullivan, I consider two kinds of interpretable neural networks: "disentangled" DNNs and approaches that combine Deep Learning with Symbolic Regression (e.g., Cranmer et al., 2020). Through these comparisons, I show that the explanatory work of these models is done by the hypotheses that they provide (in the case of neural networks) or by the hypotheses that they are in part based on (in the case of non-ML models). To make this clearer, I draw on a classification of scientific hypotheses –based on their explanatory power– made by Bunge (1997). Namely, black box, gray box, and translucent box hypotheses. I argue that while interpretable DNNs can provide black box hypotheses (i.e., those that answer questions of the "What is it?" type), opaque DNNs cannot. I show that this contrasts with the built-in gray/translucent box hypotheses (i.e., those that answer questions of the "How does it work?" type) of mechanistic models. Finally, I discuss the potential use of explainability techniques (not to be confused with interpretable ML) to extract black box hypotheses from opaque DNNs.

Luchetti, Michele & De Benedetto, Matteo: "A dynamic model of theory choice: epistemic values as environmental niches"

Kuhn (Kuhn, 1977) famously argued that scientists cannot rely on a universal algorithm to choose amongst rival theories. This is the case because the criteria of theory choice (e.g. empirical adequacy, simplicity, fruitfulness, etc.) may conflict with one another in dictating which scientific theory to choose. This conflict precludes the possibility of a universally valid choice rule. Okasha (Okasha, 2011) recently showed how the formalism of social choice theory can be re-interpreted to clarify this issue. More specifically, he claims that theory choice criteria can be seen as formally equivalent to single voters in social choice theory, while scientific theories play the role of possible choices. Based on this analogy, he argues that a famous problem in social choice theory, i.e. Arrow's impossibility theorem, holds for theory choice, too. Okasha's analogy has stimulated a lively discussion, leading to the formulation of several escape routes from this impossibility result. – In this talk, we point to an aspect of scientific theory choice that has not been explored by this debate, namely, the fact that epistemic values, just like scientific theories, are historical entities. As such, these values are influenced by the cumulative impact of previous choices. More exactly, the weight of a given value in a specific theory choice context partly depends on the outcome of previous related choices. As Kuhn (Kuhn, 1990) himself remarked, epistemic values and scientific theories evolve co-dependently in a relationship analogous to the one between environments and organisms modeled by niche-construction theorists. In other words, epistemic values have a selective role in the process of theory choice, but they are also indirectly affected by this selection process via a feedback-loop mechanism. For instance, repeatedly choosing highly accurate theories increases the weight of accuracy as a criterion for theory choice in neighboring epistemic contexts. Therefore, we claim that a more realistic account of theory choice must take this feedback-loop effect into consideration. – We model the feedback-loop effect that past choices have on the weight of epistemic values within a dynamic framework of theory choice. In this framework, the formal machinery of social choice theory is augmented with weights associated to voters and a parameter representing the time-point at which a given choice takes place. We show four possible characterizations of the impact that previous choices have on the weight of epistemic values in subsequent choices. These four characterizations formalize four distinct subclasses of the aforementioned feedback-loop effect, differing with one another as to the polarity and the strength of this effect. Finally, we show how this dynamic model of theory choice can adequately account for the mutual influence of subsequent theory choices and epistemic values in a specific historical example: the controversy between Mendelians and biometricians in early 20th century biology. This example shows how our dynamic model gives a less idealized picture of scientific theory choice than its static predecessors.

Lund, Matthew: “Bessel and the Epistemology of Observational Relativity”

In 1796, Astronomer Royal Nevil Maskelyne fired his assistant, David Kinnebrook, for “observing the times of the Transits too late.” In 1823, after analyzing the Royal Observatory's data, F.W. Bessel published a startling conclusion: ‘involuntary constant differences’ mark out the recorded astronomical transit times of distinct observers. Over the rest of the 19th century, Bessel's thesis gradually led to the framing of the ‘personal equation’ for observers, and spurred psychological investigations into the processes of visual perception. According to the dominant narrative (Boring 1929), empirical psychology's development of the personal equation put observational astronomy back on an objective footing. However, this paper argues that practical innovations within observational astronomy itself led to the stabilization of data, and that robust psychological accounts of involuntary perceptual differences only emerged after this had occurred. – This paper investigates the reasoning process Bessel, and his correspondents, went through in the years leading up to the discovery of constant differences and asks whether such apparent perceptual relativity was then viewed as a threat to objectivity. Most importantly, this paper assesses the practical reforms in observational astronomy from an epistemological perspective. – The contemporary reactions to Bessel's report were tepid and sparse, and changes in practice were put forward discreetly, without explicit reference to the problem of constant differences. Nonetheless, Bessel's discovery revealed that the epistemic terrain of astronomy was much more unpredictable than had been previously thought. Yet the solution to these worries was not a rigorous

theory of the observer, but rather a set of cautionary practices in data recording. Christoff Hoffmann characterizes the discovery of constant differences as bringing into being a ‘cold tradition’, wherein epistemic problems are “preserved in the form of undiscussed practices.” (2007, 356) – Was there a sense in which Bessel’s reaction to the phenomenon of individual constant differences was epistemically irresponsible? One might argue that individual differences degraded the data enough that their utility as guides for navigation and clock determinations should have been cast in doubt. However, the recognition of constant, involuntary individual differences really did not compromise those activities at all since their rise to levels of previously unimagined precision was due to observational data of exactly this kind. What Bessel’s revelation disclosed was a new type of error, but one whose overall effect was too small to lead to navigational mistakes or insuperable timing difficulties. Astronomical data and practice were precious commodities. Bessel was not concerned to provide a complete picture of observational psychology. He only wanted to supplement astronomy with a minimal epistemic account of observation so that the historical practices of astronomy could be preserved and extended. The ‘observer as instrument’ silently enters the picture with Bessel, as Hoffmann says, but the observations are the item of interest, not the observer. In general, this paper argues that forms of scientific practice can act as structures of epistemic support, even in advance of a tenable (and conscious) epistemology of observation.

M

Malatesti, Luca & Jurjako, Marko & Brazil, Inti: “Integrating legal categories with biocognitive data: the case of the insanity defence”

Advancements in the neurocognitive science of individuals with antisocial personality disorder might be of great significance for the application of insanity and similar defences in Law (Malatesti & McMillan 2010). However, thus far, the investigation of whether psychopathic offenders or other individuals with antisocial personality disorders should be exculpated has reached stumbling blocks (Jalava & Griffiths, 2017; Jurjako & Malatesti, 2018). – In this paper, to overcome these difficulties, we motivate and frame a proposal for a biocognition-informed recategorization of antisocial personality disorders aimed at differentiating accountable from not accountable offenders. We argue that we should not use syndrome-based categories for this task. These categories, as those used in the Diagnostic Statistical Manual of Mental Disorders (DSM), are based on observed behaviour and inferences about unobservable characteristics and personality traits without aetiological considerations. The categories of this kind for antisocial personality disorder are problematic. In fact, they have delivered small or no advances in treatment, they have low validity, they cover heterogeneous groups of people, include comorbidity, with low prospects of integration with neuroscience, genetics, and neuropsychological paradigms (Lilienfeld, 2014). In our paper, we offer instead a Research domain criteria (RDoC) type of approach for the legal case and discuss some of the conceptual problems it must address. In psychiatry, RDoC research aims at providing valid measures of disorders by integrating the data on the genetic, neural, cognitive, and affective systems underlying psychiatric conditions (Lilienfeld, 2014). The goal is to use bottom-up neurobiological data to rebuild psychiatric categories (see Brazil et al., 2018). Similarly, we propose to rethink bottom-up certain legal categories, that are needed for insanity or similar types of defences, for a more effective use of available neuroscientific data and further research. Such an approach does not need to be reductive (Jurjako, Malatesti, Brazil 2020). However, it must strike a balance between the legal categories that are determined top-down by cognitive-behavioural classifications, and specific legal normative constraints, and neurocognitive bottom-up revisions of them. Specifically, our proposal needs to address the difficult interface problem of relating folk/psychological notions and explanations embedded in the legal formulation of defences and the data offered by neurocognitive science (Francken & Slors 2018).

Martens, Niels: “Comparing the explanatory power of Λ CDM & modified gravity”

Applying the standard laws of gravity to the luminous matter observed on the night sky fails to correctly predict the evolution of that matter. Discrepancies with observations appear at many different scales: at the cosmological scale, in galaxy clusters and in individual galaxies. If standard gravity plus luminous matter gives the wrong answer, there must either be non-luminous matter, i.e. dark matter, or the laws of gravity must be modified, or both. Dark matter (Λ CDM) has been heralded as the clear winner at the scales of cosmology and galaxy clusters. Modified gravity (MG) earns its spurs at the level of individual galaxies. Dark matter simulations of structure formation still suffer from several well-known “small-scale problems” but are making progress in fitting the empirical correlations within and between individual galaxies which were once only accounted for by MG. It is not uncommon for advocates of each camp to claim that their research programme is, without a doubt, on the winning side, and that the other research programme deserves little to no attention from researchers. How can it be that each camp draws such radically different conclusions from the same data? Could it be that each camp has different (implicit) assumptions about what counts as a proper explanation of the types of empirical phenomena at hand? If so, this may deflate to some extent the tension between the mutually inconsistent victory claims by each camp: one camp may indeed provide a more proper explanation of (a preferred subset of) the data as evaluated against their (philosophical) account of explanation, whilst the other camp provides a better explanation of (a potentially different subset of) the data when evaluated against their preferred (philosophical) model of explanation. We make the case that Λ CDM advocates often wield notions of explanation that somewhat resemble unification, and that MG advocates emphasise notions that focus on a lack of fine-tuning. The first aim of this paper is to make these implicit understandings of explanation explicit, and to argue that they are, within each research programme, similar enough to indeed fall under the umbrella terms of unification and lack-of-fine-tuning respectively, even though not all aspects falling under these umbrella terms fit well with any single existing philosophical account of explanation. The second aim of this paper is to evaluate a) the strength of the model of explanation adhered to by each programme and b) the strength of each candidate explanation offered by each research programme, according to the intra-programme standard of explanation as well as the standard of explanation of the other programme. We argue that there is no programme-independent way of privileging one of the two notions of explanation over the other. It is rather the case that both are important: the programme that scores best against both standards of explanation cumulatively is therefore explanatorily privileged. We will furthermore defend the claim that each programme does not excel as much as usually proclaimed, when evaluated against their own standard of explanation. However, surprisingly, the modified gravity research programme does not do badly when evaluated against Λ CDM's unificatory understanding of explanation. The results of this paper indicate that despite the empirical stalemate between both research programmes, progress can be made by including philosophical analysis of the explanatory strengths of each programme. Not only has this pressing and large debate in modern physics received almost no attention from philosophers of science, philosophical accounts of explanation have not been tested against this modern case study (but often use old and/or artificial case studies). Finally, it is worth noting that our evaluation of the different ways in which spacetime/gravitational structures and (dark) matter fields explain may help with theory development. One example: If an explanatory weakness of one programme happens to be an explanatory strength of the opposing programme, this may provide extra motivation to take that demand for explanation seriously, as well as a guide towards developing such an explanation within one's own preferred research programme. Another example: there is a recent trend of hybrid theories, such as superfluid ‘dark matter’ theory (eg. Berezhiani & Khoury 2015, 2016), that do not obviously fall neatly within a single one of the traditional two research programmes. Understanding the explanatory strengths and weaknesses of each traditional research programme may guide development of hybrid theories that are the (explanatory) best of both worlds.

Maziarz, Mariusz: “A Perspectival View on Inconsistent Results of Clinical Trials”

Clinical trials often report results conflicting with previous outcomes (Ioannidis 2005a; 2005b; Broadbent 2013). The label ‘inconsistent results’ refers to the situation where one study delivers evidence for a positive relationship between two phenomena (variables) or a positive treatment effect, and another study suggests the two phenomena to be unrelated or even a negative sign of that relationship. The presence of inconsistent results has a detrimental impact on theoretical discourse, policy guidance, and clinical practice. Ioannidis (2005a) admitted that “such disagreements may upset clinical practice and acquire publicity in both scientific circles and in the lay press”. My purpose is to argue that inconsistent results emerge from different but plausible study designs and statistical techniques and are relevant for different clinical contexts. – Morrison (2011) and Massimi (2018) addressed the question of whether inconsistent (theoretical) models can be compromised with perspectival realism. While the pluralist view that, at least in some cases, “no one [of theories or results] is more correct than the others” (Kellert et al. 2006, xii) is present in the philosophical literature, most pluralist and perspectival positions have been established concerning models of phenomena or theories (e.g., Kitcher 2003; Mitchell 2002; Dupré 1993). The pluralism of statistical results has received limited attention. I analyze the design of clinical trials and statistical methods to argue, using case studies, that alternative methodological commitments producing conflicting results are often plausible in the sense that one cannot appraise inconsistent studies based on their methodological quality. I show that different research designs and/or statistical techniques deliver results that are representative for treatment effects in different clinical contexts or different subpopulations. This conclusion is in agreement with the position of perspectival realism. This lends support to the claim that inconsistent results are relevant for different clinical decisions. – Designing a clinical trial requires specifying an intervention, defining a control group, measured outcomes, choosing sample size, and formulating inclusion and exclusion criteria. All these choices not only can but, sometimes, do shape results. For example, the disagreement between treatment effects estimated by Yusuf et al. (2000) and Stephens et al. (1996) can be ascribed to a different approach to measuring outcome and a higher dose. The results of the RECORD study (Home et al. 2007) assessing the cardiovascular safety of rosiglitazone contrast with the finding of Nissen and Wolski (2007) due to differences in sample size, outcome (all cardiovascular attacks vs. coronary heart failure), and inclusion criteria. Furthermore, the malleability of statistical methods (see Stegenga 2018), which includes decisions regarding data preprocessing techniques, confounding factors, and testing procedures, can make two results derived from the same dataset differ substantially. These cases show that one cannot, in my view, convincingly argue that some of these choices are objectively superior to others and deliver a more reliable estimate of treatment outcome. All in all, the analysis of case studies shows that each of the conflicting results represent the biomedical reality from different perspectives and are relevant for different clinical decisions.

Meier, Lukas J.: “Thought Experimentation as a Scientific Method”

Thought experiments are mental test scenarios that purport to deliver scientifically acceptable results in the absence of actual physical execution. Scientists use imaginary situations as a method to test hypotheses or to explore the scope of concepts when the respective domain is inaccessible to ordinary experimentation or because conducting a physical experiment would be too costly, ethically impermissible, or even deemed unnecessary. Hypothetical reasoning is employed in a variety of disciplines, including in physics and economics, and it has a particularly long and important tradition in philosophical discourse, which began as early as in pre-Socratic times. The thought-experimental method has had many prominent advocates, including such major figures as Descartes and Leibniz. Appealing to intuitions about imaginary cases has also seemed dubious to some, however. A thought

experiment is, as Ulrich Kühne remarks, an experiment of which the main part is missing. And Bernard Williams worried that it is often the way in which an author describes a certain situation that determines whether or not it appears to support a particular theory, while a slightly different account of the same setting could yield entirely different results. Is this criticism well-founded? Especially in debates about personal identity, philosophers have always relied heavily on thought experiments, and the intuitions that they elicit serve as weighty evidence in favour or against the proposed accounts and concepts. Pioneered by Locke's case of the prince whose soul enters a cobbler's body and his thought experiment featuring the rational parrot, authors have made frequent use of a great variety of hypothetical situations to prove or disprove their respective views. We are, for instance, invited to envisage being teletransported to Mars or existing as mere brains in a vats. Such thought experiments are certainly ingenious and creative; but are they also suited to decide philosophical questions? – What place thought-experimental techniques should have in science in general and in philosophy in particular, as well as what epistemic status one can grant the results that this method delivers, is a difficult but very important issue. In this talk, I shall describe what I take to be the two most severe weaknesses of hypothetical experimentation. I will be arguing that especially questions of personal identity have engendered imagined scenarios that are so distant from the actual world that many of them do not comply with the standards that guide experimental design in nearly all other disciplines: objectivity, reliability, and validity. It is not surprising, then, that these hypothetical situations have been conspicuously ineffective at resolving our conflicting intuitions. I shall also show how empirically unwarranted background assumptions about human physiology render some of the hypothetical scenarios employed in the debate about personal identity highly misleading. I will illustrate each claim with well-known examples from the literature.

Meincke, Anne Sophie: "Free Will and the Metaphysics of Agency"

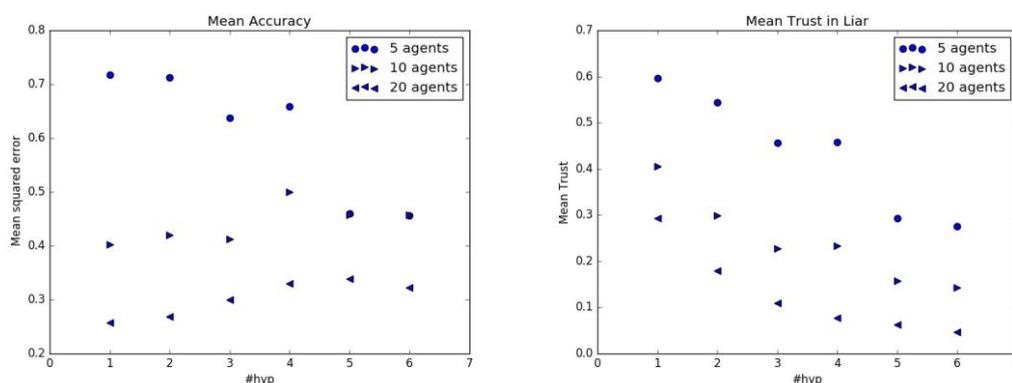
In everyday life, we take ourselves to be agents who could have decided and, hence, acted differently from how they actually did. However, within contemporary philosophy of action, free will in this robust sense is mostly doubted as it is believed to clash with the fact that humans are part of nature. So-called libertarians struggle hard to defend the existence of free will. Their appeal to indeterminism is contentious and threatens to turn our actions into a matter of chance. It seems that any attempt to naturalise free will inevitably results in eliminating what we sought to explain. In my paper I argue that the dilemma of free will is preceded by a dilemma of agency which has been overlooked as a result of the debate's traditional focus on the implications of determinism (the truth of which has mostly been taken for granted), and that we cannot hope to solve the dilemma of free will without first solving the dilemma of agency. To this end we need to rethink the ontology of agents and actions: the ultimate roots of the twofold dilemma of agency and free will lie in the metaphysics of agency – or in metaphysics more generally and fundamentally. I proceed in three steps: First, I present the dilemma of agency which is manifest in the antagonism between 'event-causal' and 'agent-causal' theories of action causation: the former's naturalist explanation in effect eliminates agency, while the latter fail to provide any explanation at all. Second, I offer a critical analysis of the underlying ontological commitments. Interestingly, both event-causal and agent-causal accounts seem to share two fundamental assumptions: (i) 'ordinary' causation is physical event causation, and (ii) all change is extrinsic to something unchanging, i.e., extrinsic to 'things' in the technical meaning of entities for whose identity change is not essential. I argue that this combination of metaphysical assumptions – physicalist mechanism and thing ontology – prevents from the outset a satisfactory account of action causation. In response, I, third, propose a novel approach which reflects and brings together two recent developments in the philosophy of biology: the study of bio-agency, i.e., the capability of organisms to interact with the environment in an adaptive manner, and the emerging turn towards an understanding of organisms as processes rather than as substances or things. My core thesis is that, ontologically, not only actions (as Helen Steward has already suggested) but also agents, qua bio-agents, are processes, i.e., entities for the identity of which change is essential. This facilitates the insight that actions in the sense of sensorimotor behaviour are just a particular, namely more sophisticated form of the interactions with the environment every

organism has to perform in order to survive. In bio-agents, agency and persistence are ontologically intertwined. I explain how the bio-process account of agency naturalises agency without thereby eliminating it, and I conclude by indicating how this may also help us with tackling the dilemma of free will.

Menon, Tarun – see Stegenga, Jacob

Merdes, Christoph: “Learning Source Reliability on Multiple Propositions”

Formal models of source reliability form an important part of philosophical inquiries into the epistemology of testimonial knowledge (cf. Merdes et al. (2020)). In particular, there exist multiple models which represent the exchange of reports between a multiplicity of agents; usually, however, the models are only applied to a single proposition (cf., for instance, Olsson and Vallinder (2013)) – unlike, for instance, in models of argument exchange – and are then used to analyze the impact of various factors on accuracy with respect to the proposition. – This field of inquiry has given rise to the notion of expectation-based updating, which describes a learning mechanism that updates the credence in a proposition as well as in the reliability of the source asserting it relative to its prior beliefs. As a side note, this is a property that those models inherit from their analytic Bayesian predecessors, though often, as with the model used in this paper, the agents are no longer perfect Bayesians, but for instance, are subject to order effects. – It turns out that trying to learn the reliability of one’s sources based on expectations is not an effective way of improving accuracy in many circumstances. This claim has been extrapolated to scenarios with multiple hypothesis. In this paper, an expectation-based model is used to represent a simple multi-proposition scenario. The model is a sequentialized, multiagent variant of the Bovens-Hartmann model of testimony (Bovens et al., 2003). The agents exchange reports on the truth or falsehood of a set of propositions and update their credences in those propositions as well as their level of trust in each other on the basis of these reports. The propositions are presumed to be statistically independent initially, and each agent is allowed to communicate at most once to each other on a particular proposition. – A scenario of particular interest is the case where a single agent is unreliable. Simulation results are depicted in this Figure:



The results confirm the claim that multiple propositions being in the mix does not by itself improve the accuracy of the group, even though there is substantially more information going around. However, the plot on the right hand side shows the mean level of trust in the unreliable agent, and both group size and, more importantly, number of propositions discussed, improve the assessment of the unreliable agent. – While this is but a small probe into the behavior of the model and the impact of adding propositions, it gives us a much improved idea of the conditions and measures under which expectation-based updating may be epistemically beneficial, but also where its limitations are. This model variant is also enriching

to the landscape of formal social epistemology in that its communication patterns – a very restricted set of reports – differs from most other models, allowing the representation of a different set of scenarios.

Michels, Robert – see Hirèche, Salim

Miller, Ryan: “Mereological Atomism’s Quantum Problems”

The popular metaphysical view that concrete objects are grounded in their ultimate parts is often motivated by appeals to realist interpretations of contemporary physics (Feynman et al., 2015; Fine, 1992; Pettit, 1993; Loewer, 2009). Given that appeals to small-scale physics are fundamentally quantum mechanical, this paper argues first that mereological atomism is only plausible in conjunction with Bohmianism, and second that it exacerbates Bohmianism’s existing tensions with serious Lorentz invariance. – Mereological atomism is only plausible in conjunction with Bohmianism because neither of Bohmianism’s leading realist competitors yields a decomposition of the physical world into a multiplicity of non-overlapping fundamental concrete objects. Multiplicity is an entirely emergent rather than fundamental phenomenon for Everettians who can’t rely on decoherence for such a decomposition (Wallace, 2012; Crull, 2013; *pace* Ney, 2021). None of the proposed ontological elements for GRW can play the role of multiple synchronic atomic parts, either. Mass density is a stuff rather than a set of atoms, flash families do not exist as concrete objects in spacetime, and the flashes themselves are sparse (Esfeld, 2014) and each occurs in the proper time of its flash family (Petrat & Tumulka, 2014) so they cannot serve as atoms synchronically composing a time-slice of a concrete object. Bohmian particles, on the other hand, provide a natural set of ultimate parts for atomists, but relying on a controverted rather than settled view in philosophy of physics comes at a significant cost. First, insofar as there is no worked out version of Bohmianism in relativistic quantum field theory, atomists are reading ontology off of either a highly inaccurate model or highly speculative future physics (Wallace, 2020). Second, and more seriously in my view, atomism compounds the problems Bohmians already have with “serious” or “fundamental” Lorentz invariance. – The basic tension between special relativity and pilot wave interpretations arises because the guidance equation requires a privileged reference frame inaccessible to any possible experiment (Maudlin, 2019). The problem for atomists is that extensional mereology requires composed objects to be determinate fusions, yet the physical particles which are supposed to constitute those atoms are not conserved in non-inertial reference frame transformations. Relatively accelerating experimenters will observe different numbers of particles (Unruh & Wald, 1984), yielding different fusions. An item of laboratory equipment appearing in the description of an experiment will correspond to one of these fusions of particles, but there will then be a hidden privileged fusion, corresponding to the hidden privileged reference frame. The actual laboratory equipment used is thus just as experimentally inaccessible as the actual time order of the experimental events. While the Bohmian interpretation of quantum mechanics might be attractive for atomists, atomism should not be attractive to Bohmians, since adopting atomism compounds the theory’s extant tension with relativity.

Minkin, Daniel: “Conspiracy Theories: Some Teachings from Philosophy of Science”

The Corona Crisis has shown—once more—that conspiracy theories permeated the democratic societies. There are theories about manufacturing Corona in a lab, about replacing the people of the

white “race” by Muslims, about lizards dressed up as powerful politicians, and so on. Theories like these are considered—correctly—as false and epistemically unjustified. But in the public as well as in the media there is also a more general picture of conspiracy theories. It seems that there are two basic tenets underlying this picture which go as follows:

- a) There is a clear-cut definition of “conspiracy theory”.
- b) A conspiracy theory is epistemically suspicious solely because of the fact that this theory is a conspiracy theory.

Though the assertions a) and b) are claimed by different scientific disciplines from natural sciences to philosophy, these claims depict scientific reality inadequately: Neither within nor between the different disciplines there is a consensus on the suitable definition of “conspiracy theory”. This is a serious obstacle for empirical and philosophical research, since we seem to need an unifying definition in order to generate a deep understanding of this phenomenon. But there is an even more serious difficulty concerning claim b): Outside of philosophy there is—as it seems—a research program which rules nearly the whole research on conspiracy theories. The hard core of this research program is constituted basically by claim b). For 15 years or so, on the other hand, epistemologists and philosophers of science has initiated an opposing research program: Philosophers like David Coady or Matthew Dentith claim that conspiracy theories don’t deserve the bad reputation they have today. These authors contend that there is no good reason to reject conspiracy theory categorically, since we have no good and general criterion as to when theories (conspiracy or not) are acceptable and when they are not. In the first part of a presentation I will outline the state of the art in philosophy, psychology and political science. It will become clear that there are far too many incompatible definitions to agree with claim a). I will also claim that there is no criterion that can help us to decide between these incompatible definitions. In the second part I want to examine famous attempts to justify claim b). There is a great number of researchers who imply, for example, that b) is right because conspiracy theories are unfalsifiable. Drawing on some positions from the philosophical debate about conspiracy theories as well as on basic insights from philosophy of science I will show that such attempts are doomed to failure. In the last part I want to make a proposal for a strategy for rejecting a) and b) without accepting highly bizarre theories like those mentioned above.

Mirkin, Julia: “Trust in Research on Human Germline Genome Editing”

“Trust in science” is considered as a common good, crucial in research ethics, policy making and public discussions. Crucial, in order to successfully realize the collective epistemic enterprise, since scientists in their everyday practice often rely on the knowledge produced by other researchers with different specialization and expertise. Crucial, since laypeople and scientists vice versa depend on each other. Laypeople depend on the knowledge provided by scientists when developing a personal stance on science-related issues and arriving at decisions about them. Scientists depend on tacit or explicit consent that underlies the public funding of science, which would not be given if people were distrustful of scientists. – There is broad international consensus, especially in the realm of novel, potentially hazardous research, that scientists must not only be trusted in their capacity as provider of information (Wilholt, 2012, p. 1). Before proceeding in the research process, to wait for and act according to the results of a broad societal discussion on ethical problems concerning their research and its possible results should be mandatory (Berg *et al.*, 1975; The Nuffield Council on Bioethics, 2016; National Academies of Sciences, 2017). Until November 2019 the debate on editing the human genome remained largely theoretical. The revelation of He Jiankuis editing of the DNA of two human embryos, who had indeed been born, has shifted the debate out of this mainly theoretical sphere (Cohen, 2019). It caused urgency in the attention that both the reactions of the Chinese authorities and of the respective scientific

community received. The results of the evaluation of this incident, including both already existing and still lacking human genome intervention regulations, can be regarded as a relevant indicator of whether the respective scientific community can be trusted in its capacity to respect the public interest. – In my talk, specifying what societal trust in science actually refers to in cases of potentially hazardous research, I present Irzik and Kurtumulus (2019) conception of and conditions for basic and enhanced epistemic public trust in science. Presenting what is known about He Jiankui’s experiment and the responses of the scientific community, I will subsequently focus on concrete questions concerning the potential clinical application of human genome editing, exposing and elaborating on the role of public trust in science. Systematically differentiating possible contexts of human genome editing, besides immediate treatment – editing embryos with a high risk of being born with a genetic disease – broader editing for disease-prevention or enhancement are also discussed. By applying the conditions for epistemic public trust, I aim to provide recommendations for gaining trust in the deliberative process of answering the focused upon questions, within the scientific community and beyond, including public actors.

Mohammadian, Mousa: “An Armstrongian Defense of Dispositional Monist Accounts of Laws”

Bird (2005) reveals an important problem at the heart of Armstrong’s theory of laws of nature: to explain how a law necessitates its corresponding regularity, Armstrong is committed to a vicious regress. In his very brief Reply to Bird (2005), Armstrong gestures towards a response that, as he admits, is more of a ‘speculation’ than an argument. Later, Barker and Smart (2012) argue that a very similar problem threatens Bird’s dispositional monist theory of laws of nature and he is committed to a similar vicious regress. In this paper, I construct Armstrong’s would-be argument in response to Bird. Then, I argue that his response causes more problems than it solves for his account of laws and natural properties. Finally, I show that Armstrong’s strategy to address Bird’s criticism can be used, quite ironically, to defuse Barker and Smart’s argument against Bird. – The structure of the paper is as follows. In Section 2, I provide a brief account of Armstrong’s theory of natural properties and laws of nature, together with Bird’s criticism of this theory. In Section 3, I discuss Armstrong’s very brief Reply to Bird (2005) and construct his would-be argument in response to Bird’s criticism. In Section 4, I show that Armstrong’s response results in some important inconsistencies between his account of laws and natural properties, on the one hand, and some verdicts of our best scientific theories, on the other hand. Section 5 briefly discusses Bird’s dispositional monist account of natural properties and laws of nature and Barker and Smart’s (2012) critique of Bird’s view. In Section 6, I argue that the strategy that is used, unsuccessfully, by Armstrong to provide a response to Bird’s criticism can indeed be successfully used to defuse Barker and Smart’s argument against Bird. Finally, in Section 7, I answer a possible objection against my argument.

N

Näger, Paul M.: “Evidence for interactive common causes. Resuming the Cartwright-Hausman-Woodward debate”

The causal Markov condition (CMC), which is a central principle of causal modelling, requires that conditional on a common cause the correlation between its effects vanishes (the common cause screens off the correlation). Since Salmon (1978) presented the first counterexamples, joined by van Fraassen (1980, 1982) and Cartwright (1988 and many more), there is a debate about whether there are also common causes that fail to screen off (“interactive common causes”, ICCs), violating the CMC. – The

deepest and most intense debate about ICCs up to date took place between Cartwright on the one, arguing for genuine ICCs, and Hausman and Woodward on the other side, defending the CMC against Cartwright's criticism. Since indeterminism is a necessary requirement for ICCs, the most serious candidates for ICCs refer to quantum phenomena, and the debate started with these. Unfortunately, early on in that debate, Cartwright focussed on non-quantum examples (especially her example of a chemical factory; first in Cartwright 1993), which could easily be shrugged off. What is more, Hausman and Woodward's (1999) redescription of quantum cases saving the CMC remain unchallenged. – This paper takes up this loose end of the discussion and aims to resolve the debate in favour of Cartwright's position. My argument comes in two steps: I first shows that all cases of purported quantum ICCs are cases of entanglement, which gives us a precise formal description of the best candidates for ICCs. I then analyse this quantum mechanical description (in a dynamical collapse interpretation) using the tools of causal modelling. – The analysis reveals that the collapse of entangled systems is best described as a causal model with an ICC. The entangled state being non-separable must be described as *one* variable and, by usual standards, it is a cause of each component of the product state after collapse; and the entangled state does not screen off the correlation between these components. – I also discuss systematically redescription of ICC structures, which try to avoid ICCs, including those by Hausman and Woodward that (i) the effects of the common cause should be redescribed as *one* variable or that (ii) there is a non-causal connection between the effects. However, option (i) fails because, according to the quantum mechanical description, the effects of the common cause are described by a product state, and product states have separate and distinct components. There is no reason to regard them as one variable. Against (ii), I argue that there neither is any convincing reason to assume that there is a connection between the components of the product state: When the product state emerges, the non-separable connection has ceased to exist. Also the appeal to conservation laws, as suggested by Gebharder and Retzlaff (2020), does not help here. – In sum, if a dynamical collapse interpretation of quantum mechanics is true, there is clear evidence that there are ICC structures in our world.

O

Oldofredi, Andrea: “Relational Quantum Mechanics and the PBR Theorem: A Peaceful Coexistence”

According to the principles of Relational Quantum Mechanics (RQM), the wave function is considered neither a concrete physical item dynamically evolving in spacetime, nor an object representing the absolute state of a certain quantum system (cf. Rovelli (1996, 2016)). In this context, in fact, ψ is defined as a useful computational tool encoding the information available to a particular observer about a specific system. Hence, it is generally claimed that RQM offers an epistemic view of the wave function. This perspective about the nature of the quantum state, however, seems to be at odds with a formal result obtained by Matthew Pusey, Jonathan Barrett and Terry Rudolph—known as the PBR theorem—according to which all ontic models reproducing the predictions and the statistics of the Born rule must be ψ -ontic (cf. Pusey et al. (2012)). Alternatively stated, as Leifer (2014) pointed out, such a theorem excludes the possibility that wave functions represent our knowledge of an underlying reality described by some ontic state (usually denoted λ). Therefore, by considering that ψ actually refers to observers' knowledge conforming to Rovelli's framework, one would be led to conclude that RQM is in plain contradiction with the PRB theorem. This talk aims at showing that relational quantum mechanics is not affected by the conclusions of PBR's argument; consequently, this alleged inconsistency can be dissolved. To achieve this result, I will take into account the foundations of the PBR theorem, i.e. Harrigan and Spekkens' categorization of ontological models (cf. Harrigan and Spekkens (2010)). More precisely, it will be argued that their implicit assumptions made about the nature of the ontic state are incompatible with the postulates of RQM. Indeed, conforming to this classification, λ is taken to be an

observer-independent representation of the state of a certain quantum system. However, the relational character of Rovelli's theory requires that, in order to define what ontic states are, one uses completely different criteria w.r.t. those employed by Harrigan and Spekkens. Following the metaphysical principles of RQM, indeed, *both* ψ and λ must be relational, meaning that

- λ represents quantum systems relatively to a certain observer,
- ψ stores information that a particular observer has relatively to a given system.

Thus, in this presentation I will carefully explain which assumptions RQM makes about λ , and how they diverge with those employed by Harrigan and Spekkens and utilized by PBR to obtain their result. In addition, I will ask whether it is possible to derive a PBR-type result in the context of RQM, and will answer this question in the negative. This conclusion also shows some limitations of the PBR-theorem that, to the best of my knowledge, have not been discussed in literature. In sum, it will be argued that Harrigan and Spekkens' approach does not have the necessary formal and ontological resources to be correctly applied to RQM. This fact has a remarkable implication for our discussion: given that PBR theorem relies on Harrigan and Spekkens' classification of quantum models, but the latter cannot be used to evaluate RQM, one can safely conclude that Rovelli's theory does not lie within the scope of the theorem, avoiding any formal contradiction with it. In conclusion, we can say that RQM and the PBR theorem can peacefully coexist.

Ongay de Felipe, Íñigo: "What is the role of Philosophy of Biology with regard to the Extended Evolutionary Synthesis and why should it matter"

The role of Philosophy of science with respect to scientific research has been the focus of much controversy in the recent past of the discipline. While many philosophers, ever since the heydays of the discipline with the Vienna Circle, have argued that Philosophy of Science represents an analysis of the logic of scientific construction, others, following Quine's idea that there exists a continuum between science and philosophy, have made the case for a more substantial contribution of philosophy to the very development of scientific research and practice. This paper addresses this topic in the relatively recent light of a set of new developments in the field of evolutionary biology, which have been known of as the Extended Evolutionary Synthesis (EES). While the systematic theoretical articulation of this new scientific framework is still to be worked out thoroughly, there are myriads of theoretical and evolutionary biologists that defend that the winds of change blowing in the realm of evolution are here to stay. Many also argue that these theoretical novelties require the concerted effort of the scientists and the philosophers if the many empirical and conceptual changes posed by the EES are to be addresses successfully. This paper starts off by drawing a parallel between the current situation regarding the EES and three episodes in the history of science when philosophers and scientists happened to join forces fruitfully; namely: the advent of the theory of relativity during the first decades of the 20th Century, the foundational crisis of mathematics and the establishment of the Modern Synthesis in evolutionary biology. With that comparison in mind, the author goes on to distinguish three different levels at which philosophy and theoretical biology can (and should) interact with regards to the issues involved by the ESS: first, the analysis of the relationships between the EES and the MS as well as the logic of scientific change that such relationships imply. Second, the study of the ontology of a variety of biological notions from niche-construction to phenotypic plasticity which have proved crucial for many of the angles of the ESS. Lastly, the revision of an array or traditional concepts in philosophy (from the notion of cause to the idea of individual) in the new biological focus generated by the ESS. All in all, the paper concludes that these three levels of cooperation are of equal interest to the philosopher and to the biologist. In this regard, the point will be made that at least when it comes to the EES a continuum between science and philosophy does obtain.

Ouzilou, Olivier: “Social sciences and conspiracy theorizing: the problem of collective entities”

Conspiracy theories have been analyzed as a form of pseudoscience in that they mimic the epistemic practices of science (Blancke et al. 2017). More specifically, conspiracy theories seem close to social sciences: they mostly deal with social objects and have been understood as a simplifying way of apprehending complex social phenomena (Leiser, Shemesh, 2018). What links does conspiratorial theorizing have with the methodology of the social sciences? The first critique of the scientificity of conspiracy theories and their inadequacy to social science methodology is due to Popper (1962). Popper considers that conspiracy theories are the antithesis of what social sciences should do, namely to analyze the unintentional effects of the aggregation of individual actions. He further argues that part of the social sciences is affected by conspiracy theorizing. According to Popper, the "sociological conspiracy theory" makes at least two mistakes concerning macro-social phenomena:

- (1) it claims to explain the existence and functioning of certain social realities (wars, famines, etc.) or institutions (school system, government) by conceiving them as the effects of the deliberate project of certain social agents
- (2) it considers social groups (nations, classes, etc.) as agents of a conspiracy as if they were individuals.

I will focus on (2), which has been much less analyzed than (1). What is the problem with this way of conceiving macro-social realities? One can interpret (2) in two ways, namely as:

- (a) a confusion between organized collective entities and unorganized collective entities
- (b) a functional theorizing of the behavior of non-organized collective entities.

I defend the relevance of (a): even it has been little analyzed, many conspiracy theories are indeed based on a confusion between kinds of groups. More precisely, I show that they stem from a category mistake in that they attribute collective intentionality to a set of individuals on the basis of some of their (supposed) common properties. (b) makes Popper's statement more questionable. The idea that functionalism in the social sciences mobilizes conspiracy theorizing is widespread among advocates of methodological individualism (Boudon, 1981; Elster, 2015). Is it legitimate? At first sight, this is a misunderstanding: the specificity of the functional explanation of a social reality lies in the fact that no agent is supposed to be aware of its function. However, after showing the superficial similarities between conspiracism and functionalism, I try to find, on the basis of the work of Bourdieu and Passeron (1990), the criteria that make certain functional explanations conspiratorial. Broadly speaking, these criteria are of two kinds, concerning: i) the content of the explanation (when the explanatory role of the function is actually similar to that of an intentional action and the functionality is harmful to the society); ii) the cognitive attitude underlying the explanation (when it consists in an epistemically irrational unification of unrelated elements under a teleological scheme and when functional theory contains what is called "strategic immunization" and "epistemic defense mechanism" (Boudry, Braeckman, 2011)

P

Paitlova, Jitka: “The value of value neutrality”

In the paper, I will focus on one of the most important concepts in philosophy of science, value neutrality. The appeal to the value neutrality of science can be understood as a specific consequence that is already known from Hume's law: at least for science, there is a rule that normative statements cannot be legitimately inferred from descriptive statements. The value-free ideal is still the general canon for every scientist who should be impartial, which means being aware of his/her own value judgments and avoiding letting them affect his/her research, particularly his/her scientific theories. To date, some philosophers have understood neutrality and impartiality as one of the key conditions for scientific

objectivity (Popper 1988, Albert 1991, Collier 2003), some as an expression of the essence of scientific inquiry (Lacey 1999). Nevertheless, an increasing number of philosophers have doubted the value neutrality of science or completely rejected this ideal because they claim that methodology cannot eliminate all bias or value impacts for systematic reasons—for example, underdetermination, social determinism, the argument from inductive risk, the problematic distinction between epistemic and other values, etc. (Longino 1990, Douglas 2009, Putnam 2016). This criticism sounds plausible but if philosophers reject neutrality, they should also renounce objectivity, and then science would fall into some form of (more or less strong) relativism. All forms of epistemic relativism claim that it is impossible to show in a neutral way that one epistemic system (such as science) is epistemically superior to the others (Kusch 2016). – I aim to suggest that the value neutrality of science is possible (in a specific, redefined sense) and, at the same time, the fall into relativism is not necessary. There is a wide range of arguments pro et contra value neutrality and I am convinced that a deeper understanding of these various arguments in the context of all relevant disputes (particularly in German philosophy) is indispensable because some new arguments presented by Anglo-Saxon philosophers lack a more profound theoretical anchorage in German philosophy which is reflected in the confusion of value neutrality with the autonomy of science. On the other hand, Anglo-Saxon philosophy has emphasized the problem of scientists’ responsibility which becomes obvious after rejecting value neutrality, and there is a whole host of other problems. I want to focus on the main current form of criticism – semantical (Putnam) and methodological (Douglas) – and to find a constructive starting point for a new redefinition of the value neutrality concept that allows the ideal of value-free science to continue to apply. I will argue that the appeal to the informative (not normative) character of scientific statements (Albert 1991) is crucial for value neutrality and, moreover, that the ideal of value neutrality is central to shaping the responsibility of scientists because scientists still consider value neutrality (or impartiality) to be a major part of their moral code and it is the basis of scientific ethics (Betz 2013, Bird 2014).

Perry, Eoin: “Representational Risk and the Representation of Statistical Evidence”

Increasingly popular in the social sciences since 2012 (see Bakker et al. 2020), pre-registration involves publishing a written plan, outlining in substantial detail how a study will be conducted, before collecting or observing (some portion of) the data for that study. Following the perceived ‘reproducibility crisis’ (see Fanelli, 2018), it is now viewed by some social scientists as part of a ‘revolution[ary]’ (Nosek et al., 2018), ‘Manifesto for Reproducible Science’ (Munafò et al. 2017). The recent rise of pre-registration in the social sciences has led to the popular uptake therein of a novel theory of statistical evidence, which I call ‘Error Statistical Confirmatory Testing’ (ESCT). – ESCT involves distinguishing ‘confirmatory’ hypothesis tests, which may provide evidence for hypotheses, from merely ‘exploratory’ tests, in which ‘p-values lose their [evidential] meaning due to an unknown inflation of the [relevant type-1 error rate]’ (Nosek & Lakens, 2014). ESCT is best understood in terms of: (1) the (2008) version of Mayo’s Error Statistical theory of evidence, and (2) a characterization of Mayo’s (2008) concept of the ‘relevant error’ with reference to Gelman & Loken’s influential (2013) ‘Forking Paths Problem’. – ESCT comes with a perhaps surprising agent-relativity implication. According to ESCT the same background information, data, relevant error, normatively desirable error rate, and calculated test statistics, may facilitate a confirmatory test of a hypothesis for one scientific agent, but a merely exploratory test of that same hypothesis for another agent. This is the case because two agents may, irreversibly once the data are in, have used different decision rules to accept the same hypothesis. Agent-relativity has been missed because of a blurring (see Mayo; 2018a, 2008, Staley; 2004) of the distinction between decision rules and data-generating processes. At the cost of agent-relativity, incorporating agents’ personal decision rules into the characterization of the evidence, as occurs in ESCT, is a relatively procedurally objective means towards two benefits: Firstly, by making it difficult to substantially change analytical strategy post-data, while still following a rule yielding a confirmatory test, decision rules prevent non-ideal

agents from finding biased ways to construe the data in support of a preferred claim. Using an agent-neutral evidential framework leaves an agent free to psychologically fool themselves about how much evidence for a preferred claim that framework actually indicates. Similar to informative Bayesian priors, agents' personal decision rules also leverage pre-data theoretical knowledge, reducing type-2 error rates (see Fluornoy 2021) compared to frequentist frameworks demanding stricter multiple comparison correction. – Despite this second benefit, I point to some evidence (Scheel et al. 2020) for the following empirical conjecture: relative to other currently popular evidential frameworks (including Bayesian frameworks), agents using ESCT will tend to commit fewer 'relevant type-1 errors', but more 'relevant type-2 errors'. In some ways echoing Rudner's classic (1953) inductive risk argument, but more aptly described as a case of *representational risk* (Harvard 2020) regarding the choice of evidential framework, I argue that the justification for encouraging ESCT in some domain will largely rest on a normative judgment about the relative costs of type-1 vs type-2 error.

Peruzzi, Edoardo and Cevolani, Gustavo: “Defending (de-)idealization in economic modelling: a case-study”

Theoretical models in science, and in economics in particular, typically contain idealizations of various kinds. Interestingly, while the idea of idealization is widely studied and central to the recent philosophical debate (Potochnik 2017; Niiniluoto 2018; Levy 2018; Mäki 2020), the companion notion of de-idealization has attracted much less attention. Roughly, deidealizing a theory or model means removing one of its idealized assumptions and replacing it with a new one that it is less idealized, i.e., more realistic in being closer to the actual phenomena (Nowak 1980; Cools, Hamminga, and Kuipers 1994; Niiniluoto 2002, 2012; Hindriks 2012; Knuuttila and Morgan 2019). – In recent discussion on the methodology of economics, the notion of deidealization and its role in the practice of the discipline has been strongly criticized (see, in particular, Alexandrova 2008; Alexandrova and Northcott 2009; Reiss 2012). Despite having different views of idealizations and economic modelling, such critics agree on one point: de-idealization strategies are actually not used in economic modelling, for the good reason that they are either unfeasible or useless. – This paper aims at rebutting this criticism and defend the viability of deidealization strategies in economics. We present a detailed case study from the theory of industrial organization, discussing three different models, two of which can be construed as de-idealized versions of the first. The baseline model – the so-called Bertrand model – contains, among others, two crucial idealized assumptions: perfect homogeneity of goods and perfect information among consumers. These assumptions have been gradually de-idealized by researchers and more realistic models have been built. In particular, we focus on the Bertrand model with differentiated goods (Singh and Vives 1984) and the Varian (1980) model of sales. We conclude that recent pessimism about de-idealization in economics is unfounded, and that de-idealization strategies are not only possible but also widely employed in economics.

Pfeifer, Niki: “The probabilistic turn in the psychology of reasoning: a necessary paradigm shift?”

The new paradigm psychology of reasoning is characterised by using probabilistic approaches instead of logic as a rationality framework. From the first experiment on deductive reasoning (Störring, 1908) until the early years of this century, bivalent truth-functional logic used to be the dominant rationality framework for theories of human reasoning. Consequently, rational reasoning behaviour was characterised by correct decisions about logical validity of arguments. Likewise, judgements about the truth/falsehood of conditionals were interpreted to be rational, if they are consistent with the truth-semantic of the material conditional. During the last two decades, however, more and more psychologists of reasoning turned to using probabilistic approaches instead of logical ones as the normative gold standard. This probabilistic turn has been characterised as a paradigm shift within the

psychology of reasoning. In my talk, I critically assess this paradigm shift from a philosophy of science perspective. – My talk starts by explaining the role of logic in psychological theory building—specifically in selected psychological theories of reasoning (i.e., the mental model theory and the theories of mental rules and mental logic)—and in the construction of experimental-psychological reasoning tasks (like truth table tasks and the suppression task). Then, I give reasons for the shift from logic to probability. Theoretical reasons include the insight that people tend to evaluate sentences not exclusively by true or false—rather they evaluate them by degrees of belief. Moreover, logic is monotonic, i.e., additional premises cannot invalidate a logically valid argument. Furthermore, people’s judgements about the truth conditions of conditionals are systematically inconsistent with the truth conditions of the material conditional, i.e., they violate the semantics of a key notion of logic. Data from probabilistic truth table tasks suggest that most people interpret their degrees of belief in indicative conditionals by conditional probabilities, which cannot be modeled by logic. I make a strong case for using coherence-based probability logic (CPL) as a rationality framework for the psychology of reasoning. As the name suggests, it is based on the subjective notion of probability, which was developed by de Finetti and which was later generalised to conditional probability. It combines ideas from logic (rule-based reasoning, argument forms, etc.) and probability theory (degrees of belief, nonmonotonicity, etc.). I explain why CPL is fruitful in the construction of psychological theories and tasks. CPL provides a unified approach for investigating the rationality of human inference in a variety of different tasks ranging from truth table tasks, nonmonotonic reasoning, reasoning about quantified statements and how people interpret conditionals. Moreover, CPL has been proven useful for studying how people assess argument strength and provides a new solution to the Ellsberg paradox in decision theory. I conclude my talk by arguing that the paradigm shift towards probabilistic approaches is normatively and descriptively appropriate and hence necessary for the progress of the psychology of reasoning.

Pils, Raimund: “Scientific Realism and Epistemic Risk”

I transfer recent developments from the value debate in epistemology to arrive at novel insights concerning the scientific realism debate. The focus is on two connected arguments, one on the value side and one on the evidence side. The value side: Suppose one of our best supported scientific theories shows *prima facie* an ontological commitment to electrons. Suppose further that electrons exist. In taking the realist stance and believing in the unobservables postulated by our best scientific theories, realists will believe an (interesting) truth. Now suppose that electrons do not exist. Then realists will believe a falsehood. Presupposing that our current best scientific theories will be sometimes right and sometimes wrong about the existence of unobservables, realists will believe more interesting truths but also more falsehoods than anti-realists. The question becomes: Is it worth to believe those falsehoods for acquiring those truths? This picture suggests that realists are more risk-seeking, anti-realists are more risk-averse, and selective realists try to walk the line in between. – This raises the question how much weight one should put on believing truth and how much on avoiding error. I call this the question of balancing. I will argue for some boundaries of rationality but, ultimately, we will be left with various permissible epistemic value choices. The purely epistemological point of view is concerned with instrumental rationality, i.e. whether one takes the right means to one’s epistemic ends. The ends themselves – whether one should be a more risk-averse or a more risk-tolerant epistemic agent – cannot be a question of epistemology. This implies some degree of voluntarism towards stance-choice. However, voluntarism is not the end of the debate. Instead, I suggest that this marks a shift to a pragmatic debate. Without epistemic reasons for balancing, such balancing must come from ethical or practical goals. Also, since epistemic reasons for balancing are limited, our theories of rationality should reflect an attitude of tolerance. – The evidence side: There are some stances that block themselves from any evidence. Consider, for instance, constructive empiricists. For them, neither historical evidence nor any supportive evidence for a theory could in principle compel one to believe even very few very restricted claims about unobservables even if there were a perfect scientific track-record. But then, if such anti-realists are this risk-averse, and there is also no theoretical reason for restricting one’s standards of epistemic risk to just

the region of unobservables, vast other regions of kinds of hypotheses would also suffer from reversals and skeptical considerations. Such anti-realists are no longer merely anti-realists, but rather very risk-averse general investigator in all regions, observable and unobservable. – This is not to say that the epistemic value trade-off cannot be struck so that standards are strict. Anti-realism might emerge as one correct position. But if it does, it is so because the evidence says that this is the right verdict and not because we dogmatically insist that no evidence can ever proof realism right.

Pincock, Christopher & Poznic, Michael: “What do engineers understand? The case of biological methanation”

What sort of understanding do engineers acquire through their investigations into how to best make something? This is the leading question of the present paper. We will discuss a concrete problem in engineering that is to propose solutions to the task of synthetically producing methane as a storage medium for energy. This case will be analyzed in light of recent discussions of scientific understanding. – Philosophers of science have recently emphasized the topic of understanding for some particular special sciences but mostly for science in general. It is debated whether the understanding of researchers and lay people that has a scientific content is to be explicated in terms of accounts of explanatory understanding or objectual understanding (Khalifa 2017; Dellsen 2020; Rice 2021). Building on Dellsen’s account of objectual understanding, we focus on the engineering sciences and a particular case study. In the philosophy of technology, engineering knowledge is taken to be a special kind of knowledge and there are debates about whether scientific and engineering knowledge are two distinct kinds of knowledge (cf. Houkes 2009; Kant & Kerr 2019). A prominent question is: what do engineers know? A monograph by Walter Vincenti is often referred to in debates of philosophy of technology, which uses this phrase of 'what engineers know' in its title (Vincenti 1990). Building on both debates, the question of the epistemic achievements of researchers such as engineers is here pursued in terms of understanding. – This paper develops a case study to argue for three conclusions about engineering and understanding: (i) in line with much recent work in the philosophy of science, an engineer’s understanding turns on grasping models. However, (ii) there are at least two different kinds of models: representational models that aim to accurately depict features of target systems, and also what we call “design models” that aid in the production of artifacts. We argue that the understanding of the engineer through design models should be informed by the best available representational models. This is because an engineer is more likely to achieve a good understanding of some phenomenon when their design model is appropriately integrated with good representational models. This conclusion supports Dellsen’s (2020) recent argument that objectual understanding does not require explanation and is based on the goodness of particular models. The goodness of models is cashed out in terms of accuracy and comprehensiveness according to this argument. We argue, though, that (iii) Dellsen’s account of understanding based on models of dependencies has to be amended in order to make sense of the engineer’s understanding. The goodness of a model influences the degree of understanding a subject can obtain, and this has repercussions for the engineer’s understanding that is based on grasping design models. As the goodness of a design model is determined by different criteria than the goodness of a representational model, Dellsen’s account must be extended. We will argue that the criteria for design models are (a) comprehensive and well-structured information, (b) user-friendliness and (c) overall goodness of the product.

Poznic, Michael – see Pincock, Christopher

R

Raab, Jonas: “Extended Abstract: Too Many Dutch Book Arguments?”

The purpose of this paper is to develop a new argument against so-called Dutch Book Arguments (DBAs), and to discuss two potential objections. Let me present the argument before briefly indicating the objections and my replies. The DBA intends to establish *probabilism*: The thesis that one’s degrees of belief *should* be probabilities. As the ‘should’ signifies, it’s a *normative* thesis. The DBA consist of two pieces of mathematics together with the connection to rationality. The pieces of mathematics are the Dutch Book Theorem (DBT) and its converse (CDBT). The DBT shows that if one’s degrees of belief are *not* probabilities, then one is liable to accept seemingly fair bets which lead to net loss. On the other hand, the CDBT shows that ordering one’s degrees of belief accordingly, one is secure of such net loss bets. The connection to rationality is that accepting bets which lead to net loss cannot be rational. What the two theorems imply is that having degrees of beliefs which are probabilities protect us from such irrationality. Hence, we should order our degrees of belief accordingly. So much for the standard DBA. Recently, the concept of probability has been generalized. In particular, the above DBA is targeting a notion of probability based on *classical* logic. However, already de Finetti’s (1974) proof is more general than is needed for this particular version; Paris (2001) shows that there is a general characterization of probabilities with DBTs and CDBTs. The particular characterization is not important for our purposes. What is important is that the above DBA is likewise more general; it applies to *all* probabilities which satisfy the characterization. Besides classical logic, strong Kleene (K_3) satisfies it (see Williams 2012, 2016). Also note that I did not mention the underlying logic in my description of the DBA. That means that it invariantly applies to *all* probabilities which have DBTs and CDBTs. In particular, this means that if *one* DBA is correct, *all* have to be. – This, though, is problematic. Consider the following:

(LEM) $\phi \vee \neg\phi$.

According to classical logic, (LEM) is tautologous. Therefore, $b(\text{LEM}) = 1$ —and this is a *normative* claim. But strong Kleene does not have tautologies. Therefore, it is *not* normatively committing that $b(\text{LEM}) = 1$. It follows that we *are* and *are not* normatively committed that $b(\text{LEM}) = 1$ —a contradiction. Something must have gone wrong. I conclude, that the DBA cannot establish what it intends to establish. – This brings us to potential objections to my argument. I want to discuss two. The first points out that we can resolve the issue by distinguishing between different degrees of belief. The second objection points out that not all DBAs are on a par, but that we have a tie-breaker: the argument relies on classical logic in the background so that the classical DBA can be saved. – I argue that the first objection fails because it conflicts with the motivating idea to be secure from certain loss. And I argue that the second objection fails for logics capable of proving such general results; however, choosing a weaker logic might be too weak, i.e., not constrain our degrees of beliefs adequately.

Reinhard, Franziska: “Re-Construction or Re-Invention? Experimental Research into the Origins of Life”

Prototypical historical sciences are in the business of re-constructing the past. For the most part, they do so on the basis of traces. Traces are remains of past phenomena and processes that can still be discovered and observed today. For example, palaeontologists study fossils, geologists sedimentary layers, and archaeologists perform excavations. In this talk, I will focus on a rather non-typical historical science: origins of life research. On the basis of that, I will draw a distinction between re-construction and re-invention as research methods in the historical sciences. – Origins of life researchers try to account for how, where, and why life first emerged. Origins of life research is a vast interdisciplinary endeavour. I

will focus particularly on so-called ‘prebiotic chemistry’, a subfield drawing on organic chemistry. Prebiotic chemistry seeks to understand how biomolecules (from amino acids and simple carbohydrates to nucleic acids, proteins, and lipids) formed from simple precursors; how they self-assembled and ultimately gave rise to biological functions such as replication or metabolism. By our current estimations, these processes took place on the early Earth more than 3.5 billion years ago. This makes for a particularly challenging epistemic situation even compared to other historical sciences. Researchers cannot expect to find substantial traces of the formation of first forms of life. In addition, their knowledge of environmental conditions on the early Earth is uncertain. – To counteract this difficult epistemic situation, prebiotic chemistry heavily relies on techniques and concepts from standard organic synthetic chemistry. Importantly, the research is largely experimental and focuses on synthesizing rather than analysing relevant molecules – even though the goal is to explain a long-completed historical process. – The rationale behind this approach is summarized by one OoL researcher as follows: “... biogenesis, as a problem of science, is lastly going to be a problem of synthesis. The origin of life cannot be ‘discovered’, it has to be ‘re-invented’” (Eschenmoser, 2007). – In my talk, I will make precise the notion of re-invention. In contrast to historical re-construction, re-invention does not focus on drawing inference on the basis of existing traces as the main element. Rather, it brings experimental research in the historical sciences to the forefront. I will argue that whether or not re-invention as a method of accessing the past is adequate depends on the epistemic situation at hand as well as the integration of different re-inventions with other types of evidence. Finally, I will draw out the consequences of the distinction between re-construction and re-invention for debates about the general distinction between the historical and experimental sciences.

Reydon, Thomas – see Desmond, Hugh

Rivat, Sébastien: “How Theoretical Terms Effectively Refer”

Selective realists who follow most closely the traditional form of scientific realism, such as Kitcher (1993) and Psillos (1999), usually acknowledge that the problem of referential failure across theory-change requires adjusting both one's semantic and epistemic commitments. For instance, if we grant that at least some of the central terms of our "best" past theories fail to refer to anything real, we cannot assume that, in general, the terms of successful theories automatically refer to the right sorts of entities and restrict ourselves to selecting descriptions that we can trust. We also need to account for: (i) the mechanism by which some, but not all, theoretical terms come to refer to unobservable entities; and (ii) the putative referential stability of some, but not all, theoretical terms under theory-change. – The central challenge underlying both (i) and (ii) is to find a reliable and principled way of distinguishing between referential success and failure, i.e., a principle of selective reference, and this is far from trivial. For instance, we cannot appeal to the theoretical content of our best current theories since we do not yet know whether they will not appear to be deeply mistaken by the light of future theories. Nor can we point to the crucial predictive and explanatory role of a term since the next theory might show that, ultimately, this term was not playing such a crucial role. In other words, and as Stanford (2006) has made it particularly vivid in my view, we need to find a reliable and principled way of distinguishing between referential success and failure before theory-change. – The goal of this talk is to present a theory of reference which allows selective realists to address both the traditional problem of referential failure and Stanford's challenge. I will first engage with Psillos's (1999, 2012) theory of reference and briefly show that it suffers from a pernicious type of referential indeterminacy. Then, drawing Psillos's account, I will propose a theory of reference (CST*) inspired by the paradigm of effective theories developed by physicists in the 1960-80s and argue that the principle of selective reference at work in (CST*) offers a reliable way of identifying stable referents before theory-change. In short, the main idea is to take the term t of a theory to refer if t picks out some entity specifiable within the empirical reach of the theory

and assess referential success at a given time accordingly. I will show with a simple example of Newtonian physics that (CST*) is remarkably reliable in the case of effective theories. And if time allows, I briefly conclude by explaining how (CST*) works for problematic historical cases.

Rorot, Wiktor: “Counting on the Cilia: Cybernetics, Morphological Computation, and Computational Enactivism” (Symposium “The Cybernetic Renaissance”)

The recent re-discovery of the cybernetic tradition (e.g., Seth 2015; Linson et al. 2018) emphasises its impact on both the computationalist and anti-computationalist (embodied and enactive) approaches to cognition (see Dewhurst 2019). Historically, there has been a lively debate between early proponents of computationalism and cyberneticists, to mention only the interaction between Alan Turing and British members of the movement (see e.g., Pickering 2010). Similarly, the development of the enactive approach has repeatedly invoked the cybernetic tradition (e.g., Varela 1986; Varela, Thompson, and Rosch 1991). These two perspectives, however, are traditionally regarded as antithetical. Computationalism relied on the semantic view of computation (Fodor 1975), which required symbolic and representational operations. Enactivists regard the notion of representation as incompatible with their views on the autonomy of biological systems (Dewhurst 2019). However, recently several attempts have been made to reconcile the two traditions (e.g., Villalobos and Dewhurst 2018; Korbak 2019). – In this context, the research on morphological computation becomes an important piece of the puzzle. Morphological computation is “computation obtained through interactions of physical form” (Paul 2006, 619). While there is some controversy about which processes constitute examples of morphological computation, this general definition is quite unproblematic. Müller and Hoffmann (2017) identify three distinct types of processes which are (in their view, incorrectly) subsumed under the notion of morphological computation: (1) morphology facilitating control, (2) morphology facilitating perception, and (3) morphological computation proper. This distinction results from constraining the definition of physical computation into one which requires the operation of encoding, decoding, and a user who treats the physical systems in question as a computer (Müller and Hoffmann 2017). This is radically different from the views on computation that Dupuy takes cyberneticists to hold (Dupuy 2009, 6), according to which computation is a purely mechanical operation, “devoid of meaning.” Under this view, at least “intrinsically computational” processes (Crutchfield, Ditto, and Sinha 2010; Müller and Hoffmann 2017) should expand the domain of “morphological computation proper.” However, further developments of cybernetics saw a departure from this view in the form of “second-order cybernetics” (see Froese 2010, 2011; Pickering 2010), which includes and appreciates the role of the observer, the scientist studying the system. This modification sought to accommodate, among others, the apparent impossibility of self-organisation in accordance with the principles of early cybernetic views. – The purpose of the talk is to explore consequences of the cybernetic view on morphological computation. The case study will be the processes of bioelectrical communication which have been indicated both as the evolutionary origins of neural activity (Prindle et al. 2015), and as the principle tying together the activity of multicellular biological systems (Levin 2019, 2021). I will argue that the example provided by the role of bioelectric communication in morphogenesis (i.e., the development and maintenance of complex patterns in biological systems) provides an example of morphological information processing, morphological control and morphological computation proper, and as such can be taken to support the broad “computational enactivism” project and underscore the role that cybernetics may play in this framework.

Roski, Stefan: “In Defence of Explanatory Realism”

Explanatory realism is the view that explanations work by providing information about relations of productive determination. While causation is a prime example for such a relation, other candidates such

as grounding, parthood, and realization have been discussed in recent years. Indeed, realism has gained considerable popularity especially in the context of debates about non-causal explanation (cf. Schaffer 2016; Kim 1988). What makes the view particularly attractive is that it fits nicely with the idea that not all explanations are causal whilst avoiding an implausible pluralism about explanation. Although explanations track different relations (causation, grounding, etc.), there is only one fundamental explanation-relation, and one fundamental notion of explanation picking out that relation. – In spite of its attractiveness, explanatory realism has recently been subject to criticism. Taylor (2018) has presented four types of explanations that the view allegedly cannot account for: analogical explanations, explanations by convention, explanations by *reductio ad absurdum*, and statistical explanations. – The aim of my talk is to defend explanatory realism against Taylor’s challenges. I will show that her alleged counterexamples can in fact be accounted for by realism, properly understood. In what follows I will give a brief outline of the argument of my talk. – We will understand realism as characterized by the following claim:

(ER) If some propositions Γ explain ϕ , propositions Γ provide information about entities that stand in some productive determination relation to an entity that ϕ is about.

The talk will focus on three particularly interesting types of alleged counterexamples to (ER):

- (i) explanations by conventions,
- (ii) statistical explanations, and
- (iii) explanations by *reductio ad absurdum*.

Explanations of the first type refer, for instance, to rules of football in order to explain certain events during a game (e.g. “he was sent off after a tackle because the rules of football prohibit tackling”). According to Taylor, these rules do not refer to causes or grounds of the explanandum event in question. Explanations of the second type explain certain facts by general statistical laws. An example would be an explanation of the fact that the outcome of throwing two unbiased dice sufficiently often will approximate 7 in terms of the Central Limit Theorem. According to Taylor, statistical laws like the Central Limit Theorem do not contain any information about causes or grounds of the outcome in question. Explanations of the third type explain certain facts by showing that assuming their contrary leads to contradiction. – I will show that all of these cases either satisfy (ER) or are not explanations in the first place. Explanations by convention pragmatically convey information about causes or grounds, statistical explanations rely on assumptions about causal dispositions, and Taylor’s explanations by *reductio ad absurdum* aren’t real explanations. In the background of my arguments is an account of what it means for an explanation to be *about* parts of reality that is more permissive than extant accounts without, however, trivializing (ER).

S

Salomone-Sehr, Jules & Bourgeois-Gironde, Sacha: “What might we learn about shared agency thanks to game theory?”

Many of the things we do are no more than mere individual actions: combing one’s hair, reading in bed. In addition to individual agency, however, we also possess the capacity for *shared agency*: we move heavy furniture, work in firms, and organize protests *together*. In this paper, we explore what game theory might teach us about our capacity for shared agency. It might seem, indeed, that game theory is ideally positioned to illuminate what it takes to share agency, and especially the obstacles that often stand in the way of this sharing. After all, game theory has been used to model *collective action problems*

(Olson), i.e. situations where individual agents would be better off if they acted together but where individual incentives favor defection and, therefore, encourage collective inaction. – In this paper, we pursue two aims: we caution against conceptual confusions about shared agency invited by unreflective reliance on game theory; and we identify questions about shared agency that fruitful engagement with game theory might illuminate. First, we warn against mistakes that are often made when using game theory, and especially the Prisoner’s Dilemma (PD), to model collective action problems. Perhaps because PD is usually invoked in discussions of collective action problems, the strategy profile where players receive the second highest payoff is widely conceived of as one where the players *cooperate*. This, in turn, encourages two thoughts: first, no other strategy profile in PD involves shared agency; second, all it takes to act together, in PD, is for each agent to play the ‘cooperative’ strategy. But this is misconceived. Shared agency, indeed, is a more robust phenomenon than this second thought suggests. As commonly conceived in the philosophical literature, shared agency is not a matter of individual agents acting merely in parallel and whose actions might bring about a collective outcome. Shared agency, rather, requires a genuine pooling of agency, e.g. via shared intentions (Bratman), mutual obligations (Gilbert), or the enactment of a common plan (Shapiro). Accordingly, it is false that all it takes, to act together in PD is to play the ‘cooperative’ strategy. For acting in accordance with this strategy, by itself, need involve neither shared intentions, nor mutual obligations, nor even a common plan. – This, in turn, suggests that it is false that PD’s ‘cooperative’ strategy profile is the only profile that might involve the players’ shared agency. After all, it is possible (perhaps not rational, but certainly possible) to (e.g.) share the intention to play PD’s Nash Equilibrium. – Second, we identify two questions that reflective engagement with game theory might illuminate. First, how might the common plans widely thought to be required for shared agency might have emerged? Second, what is involved not just in sharing one’s agency with others, but in doing so in a manner we would recognize as *cooperative*? We consider both questions in turn and explore how discussions of solutions concepts— e.g. Nash Equilibrium, but also we-equilibrium (Bacharach) and Kantian equilibrium (Roemer)—might shed light on cooperative forms of shared agency.

Sánchez-Dorado, Julia: “Judgments of similarity and a pragmatic account of representation”

Attention to recent debates on scientific representation reveals an increased interest among philosophers of science in advancing more pragmatic, historically and socially-sensitive accounts of how models represent natural phenomena. Weisberg (2013) has for instance recently argued that an analysis of representation should reflect judgments, particularly judgments of similarity, that scientists actually make in practice. This differentiates his account from other similarity-based approaches in which the relation of representation “holds between inaccessible, hidden features of models and targets”. Examples of this type of approach can be found in French (2003), Bartels (2006), and Contessa (2007), who propose some form of rational reconstruction of representational practices but do “not fully explored the role of theorists’ intentions in all aspects of modeling” (Weisberg 2013: 5). – In this paper, I endorse the project of advancing a genuinely pragmatic account of successful representation, and specifically a similarity-based type of account sustained on a systematic examination of judgments of similarity employed by epistemic agents in modelling practices. However, I argue that Weisberg’s (2013) own proposal, the ‘weighted feature-matching account of similarity’, is although well motivated unable to accommodate actual judgments made by scientists in practice, and therefore provide as it stands a satisfactory pragmatic account of representation. In response, I propose a more precise characterization of the notion of ‘judgments of similarity’, specifying their role and importance in a normative account of representation. – More specifically, I argue that a study of judgments of similarity can function as the source of normativity of an account of representation since agents stabilize uses and define norms concerning similarity within practices of representing. I examine historical reports of the design, construction, and evaluation of the San Francisco Bay Model (SFBM), case that Weisberg extensively discusses, in order to expose how certain judgments of similarity become entrenched and eventually

give rise to rules and standards in an epistemic community given their usefulness in securing the epistemic success of the practice (USACE 1963; Huggins & Schultz 1967). Judgments are not mere opinions, but abilities that at least in part are regulated by rules, even if they are fallible and open to revision (Brown 1988, Elgin 1996). Thus, the study of judgments of similarity can help reconcile in a single account the descriptive and the normative components of a pragmatic, socially-sensitive account of representation based on similarity.

Santos-Sousa, Mario: “Progress in Psychiatry”

In psychiatry, as in other medical disciplines, classification is paramount to guiding diagnosis and effective treatment. But psychiatry seems to be lagging behind the rest of medicine in its efforts to develop a common framework for classifying mental disorders. The fifth revision of the Diagnostic and Statistical Manual of Mental Disorders (DSM-5), though much anticipated, largely ended up in disappointment, spurring a number of alternative initiatives, such as the Research Domain Criteria (RDoC) framework or the Hierarchical Taxonomy of Psychopathology (HiTOP) framework. – Given the current lack of consensus and of shared sense of direction, it is natural to ask whether (and, if so, in what sense) there is progress in psychiatry. This question has been subject to intense debate in the years leading up to and following the publication of DSM-5, which still remains the standard approach to psychiatric diagnosis, along with its counterpart, the eleventh revision of the International Classification of Diseases (ICD-11). – These debates, however, have been carried out under the implicit assumption that the ultimate goal of any system of classification is to carve nature at its joints. As a result, monists carry on in the hope that, sooner or later, their favoured taxonomy will be vindicated, while pluralists trust that all the different pieces will eventually fall into place. Sceptics, on the other hand, will smirk at the others convinced that there is no progress in psychiatry—and none to be expected (since there are no joints to be carved). – Be it as it may, I plan to pursue a more modest approach and assess the relative merits of competing taxonomies, rather than adjudicate between them. This involves shifting the focus from measuring progress globally, against an elusive end goal, to measuring it locally, against more specific aims and values. In particular, I will pitch the aforementioned attempts at psychiatric classification against each other in light of their clinical utility and empirical validity. As things stand, there is a trade-off between the two, with DSM-5 and ICD-11 scoring higher on the former (but not the latter) and RDoC and HiTOP scoring higher on the latter (but not the former).

Sarikaya, Deniz – see Blessenohl, Simon

Sarisoy, Johanna: “A failure to replicate - a failure of what?” (Symposium “The Replication Crisis and Philosophy of Science”)

Replication has been characterized as a critical test of objectivity or a mark of the scientific. In 2015, The Open Science Collaboration conducted a meta-analysis on the replicability of results in psychological research, which has been cited more than 5000 times (2015). The report is often used to attest that psychology is in a “replication crisis”. Psychologists claim that it shows that psychology has fundamental problems. – But do the results of the Open Science project give us reason to assume that psychological research does not pass the mark of being scientific? I argue that they do not. I aim to explain how the results of the Open Science Collaboration can be properly interpreted. I argue that they give us reason to be concerned over a high rate of published false-positive claims, but they are no cause for any deeper concerns about the scientific standard of the discipline. The most commonly applied marker of replication success in psychology is what I call the significance match test SMT. In the Open Science sample of 100 replications, the significance of the p-value in the original and the replication

matched in 39 of the cases. Drawing on the Open Science SMT, Diener and Biswas-Diener claim that many results in psychological research are incorrect. I slightly disagree. I explain that mismatch doesn't necessarily indicate that results are incorrect. It is to be expected to some degree due to the variability of the p-value. Further, I argue the results of the Open Science SMT, cannot tell us how adequate (or inadequate) hypotheses or measurement assumptions are in psychological research. Under significance testing, the hypothesis and the assumptions are never tested. They are simply assumed to hold true. However, I agree with Diener and Biswas-Diener (2016) that the results of the SMT (39%) are reason for concern. They are lower than expected. I argue the results point towards a high rate of published false-positive results in the Open Science sample. I argue that the drive towards publishing novel and exciting results - which reinforces publication bias and questionable research practices - leads to a high rate of falsepositives in published research and explains the low significance match rate. This, however, does not mean that psychological research is not objective. In fact, psychology is one amongst many disciplines with a high false-positive rate due to publication pressures. Several solutions such as open science and registered reports have been proposed. I conclude, although we should be concerned over falsepositives rates in published literature, a low significance match rate is no reason to believe that the discipline does not pass the mark of being scientific. If we want to know more about how well psychology is doing, we need to refrain from understanding replication in terms of SMT.

Schrenk, Markus: "Which Predicates, which Properties for Better Best Systems?"

Many advocates of the Better Best System Account (BBSA) of laws of nature suggest that Lewis-style best system competitions (BSA) can be executed for any arbitrary but fixed set of predicates/properties. This affords the possibility to launch system analyses separately for each of the special sciences (e.g. Cohen & Callender 2009, Author 2008). – However, predicates/properties of these sciences can cause trouble. In Lewis's original best system analysis, predicates refer to perfectly natural properties only, i.e. we have canonical language-to-world fit. This possibility of a smooth transition is taken for granted in most formulations of the best system idea. Yet, when turning from Lewis's BSA and his natural properties to the BBSA, the transition from the BBSA's (non-natural, scientific) predicates to the respective properties is not straightforward. Indeed, it is surprisingly hard to find a semantics for the predicates of the special sciences that suits the purposes of the BBSA. – In this paper, I will consider (i) semantic externalism, (ii) reference magnetism, and (iii) description theories of reference and will find all of these theories wanting for the purpose of the BBSA: The first tends to run against the Humeanism that is at the core of BBSAs, and the second is prone to collapse into Lewis's BSA. – While the third option is the most promising, it puts the BBSA in danger of being circular: in a description theory of reference it's the predicates' intensions, notably causal/nomological roles, that belong to their semantic content and that fix their extensions. Yet, if this reference-fixing mechanism chooses the predicates' extensions already for the nomological roles they fulfil then the BBSA seems to be obsolete: the BBSA was meant to deliver the nomological facts, yet, they are already present prior to the BBSA. I propose three answers to this problem:

(1) While the intensional roles attached to predicates discern their extensions, it is only these extensions, i.e. the 'naked' properties void of such roles, that are systematised in BBSAs. I.e., the intensional roles are mere reference-fixers. The scientific predicates' intensions can be treated as epistemic agents' nomological conjectures. Not these 'hypotheses' count as the 'real' laws, only those nomological roles the BBSAs will deliver do.

(2) Relatedly, as scientific progress shows, some of the predicates' intensions (nomological role hypotheses) might well be wrong: scientists will probably err. That such discrepancies between prior conjectured roles and anterior ideal BBSA outcomes are likely diminishes the danger of predetermined outcomes.

(3) The predicates' intensions are, finally, most probably not exhaustive. The best system might well list some additional axioms or theorems involving global matters which are not yet captured locally by the predicates' prior intensions.

Still, the reference fixing roles of the predicates' intensions do at the very least introduce a bias into the mosaic of objects the BBSA is supposed to systematise, however small it might be. For BBSAs, an innocently given mosaic is a myth.

Schroeren, David: "State-Space Fundamentalism: An Escape from the Pessimistic Meta-Induction"

The pessimistic meta-induction (PMI) is an argumentative strategy aimed at undermining a central tenet of scientific realism: that we are justified in believing that the best theories of current natural science are (approximately) true. Roughly: (D) if past theoretical transitions overturned the ontological posits of previous theories, for all we know so will the transitions from present to future theories; (H) we are justified in believing the ontological posits of present theories only if we are justified in believing that they will be preserved under future theoretical transitions; (I) but past theoretical transitions have overturned ontological posits, so (J) we are not justified in believing the structural posits of present theories. – The goal of this paper is to assess the extent to which state-space fundamentalism—a radical proposal for the ontology of modern physics recently brought into focus in the literature—is might be capable of escaping the PMI. State-space fundamentalism is a metaphysical doctrine inspired by one of the central paradigm shifts in physics during the HMth century. During this time, physicists noticed the incredible predictive and explanatory power of theories that give center stage to the notion of a state space. The state of the world at an instant of time is a maximally specific and exhaustive proposition, in that it entails the value of every physical quantity at that instant. Within the old Newtonian paradigm, state spaces are useful but ultimately ontologically dispensable: facts about the state of the world are determined entirely by the properties of the physical objects that make up the world. But the quantum revolution forced a profound transformation of this received wisdom: within quantum theories, physical states are not summaries of more fundamental goings-on whose physical significance is derived from the basic entities whose properties they describe; rather, states are promoted to the fundamental entities in their own right. This paradigm shift amounted to nothing less than a revolution: just on the basis of claims about the state space of the world, physicists were able to make inferences about the nature of physical properties as well as about the nature and behavior of matter—claims that were of unprecedented empirical accuracy. This paradigm shift has been taken as a motivation for state-space fundamentalism: the thesis that the world is a point moving in fundamental state space, whereas the familiar picture of the world as composed of physical objects located in three-dimensional space is non-fundamental and metaphysically derivative on the world's location in state space. – As this paper argues, state spaces play an irreducible and central role not just in present and past quantum theories, but also in every major proposal for unified quantum theories of gravity and matter, including loop quantum gravity and string theory. While this does not amount to a guarantee that future physical theorizing will preserve the centrality of state spaces in current physics, this paper argues that the metaphysical framework afforded by state-space fundamentalism nonetheless has a good claim to be applicable *mutatis mutandis* to future physics and thus to withstand the PMI by resisting premise (D).

Sekatskaya, Maria: "Reductionism in the Philosophy of Science and the Problem of Mental Properties"

Reduction in the philosophy of mind is usually understood in a very strong sense: as a complete reduction of all mental predicates to physical predicates (Fodor 1982; Kim 1993). In the early stages of logical empiricism, this type of reduction was considered to be about explicit definability/translatability

of theoretical predicates with the help of empirical predicates. Typically, in philosophy of mind, the accounts that do not subscribe to this type of reduction of mental concepts are classified as non-reductive accounts (Clapp 2001; Walter 2006). This gives the impression that all non-reductive accounts have something in common. In particular, non-reductive physicalists often claim that mental phenomena have a special epistemological status and therefore differ significantly from other natural phenomena. This claim is then used to justify the postulation of differences in ontology. If mental predicates cannot be explicitly defined in terms of physical predicates, then mental properties cannot be reduced to physical properties. However, the step from the failure of explicit definability of mental concepts in terms of physical concepts to proclaiming that mental phenomena are ontologically non-identical to anything physical does not appreciate the complexity of different forms of scientific reduction. In philosophy of science, explicit definability is considered the strongest, but not the only possible, form of reduction. A weaker form of reduction is that of employing bilateral reduction sentences for theoretical predicates such as dispositional terms (cf. Carnap 1936/37). But even this approach was quickly found to be untenable, for which reason a weaker constraint of reduction in terms of empirical confirmability of propositions with theoretical predicates was put forward in the classical empiricist programme (cf. Carnap 1950/62). – Although, historically speaking, logical empiricists such as Carnap and Herbert Feigl took the case of psychological theorizing as a paradigm case for discussing scientific reductions, it seems that the discussions in the philosophy of science and the philosophy of mind have diverged quite a bit and lost relevant points of interaction. In this talk, we outline a framework for better interrelating the discussions. We propose a mapping of different accounts in the philosophy of mind based on the three types of scientific reduction explained above. We argue that eliminativism, particularly type- and token identity theories of the mental, are versions of reductions in the sense of explicit definability, whereas functionalism can be framed as a form of reduction by the help of bilateral reduction sentences: functional definitions of the mental are coarse-grained, similarly to dispositions in the bilateral reductive accounts in the philosophy of science. The latter fact is not very surprising: historically, early dispositionalists can also be seen as both functionalists and physicalists (Ryle 1949; Smart 1959); the controversy between functionalism and reductive physicalism arises only at a later stage (cf.: Block 1978), and is argued against in contemporary approaches (Clapp 2001). Our grouping together of eliminativism, type identity, and token identity theories as three different versions of reduction as explicit definability is presumably more surprising, since type- and token identity theorists are realists, and eliminativists are anti-realists about the mental. We will argue that their respective realism or anti-realism comes not from the different form of reduction employed, but from a different interpretation of ontological consequences of explicit definability. Finally, we tentatively argue that supervenience accounts of the mental can be framed as either accounts of explicit definability or as accounts of reduction by empirical confirmability.

Šešelja, Dunja – see Herfeld, Catherine

Shang, Yafeng: “Is evidence of mechanisms sufficient for making within-case causal claims?” (Symposium “Mechanisms in the Cognitive and Social Sciences”)

Political scientists are interested in studying causes of rare events. What are the causes of World War I? What are the causes of the terrorist attack on the World Trade Center on 11 September 2001? What are the causal factors of the weak American welfare state? A standard method used to investigate these problems is process-tracing, which is typically defined as a method to unpack causal mechanisms (Beach and Pedersen 2013; Crasnow 2017). Many political scientists contend that it is sufficient to establish a causal claim by identifying an underlying mechanism. However, such a view is incompatible with Evidential Pluralism, which maintains that in order to establish a causal claim, one normally needs both

evidence of correlation and evidence of mechanisms (Russo and Williamson 2007). This paper defends the application of Evidential Pluralism in the context of political science by arguing that it is not sufficient to make a within-case causal claim with evidence of mechanism alone. – I begin addressing two arguments for the view that evidence of correlation is not necessary for making within-case causal claims in political science. The first argument stems from a concern that evidence of correlation is difficult to obtain in the cases of rare events. The second argument is from the observation that qualitative political scientists are not concerned with quantitative methods. I argue that neither of the arguments is compelling by showing that both arguments assume some misunderstanding of evidence of correlation. I illustrate my point with a case study of Weinstein’s work on violence in civil war (2007).

Sikimić, Vlasta: “Efficient Team Structures in Biology”

Agent-based models have been typically used to investigate the reliability and speed of the scientific pursuit in different group structures (Grim 2009, Kummerfeld & Zollman 2015, Zollman 2007, 2010) or the division of cognitive labor in science (Alexander et al. 2015, De Langhe 2014, Weisberg & Muldoon 2009). In our research, we were concerned with the optimal internal structure of research teams. This aspect has been investigated in management studies, e.g., (Gist 1987, Rulke & Galaskiewicz 2000). For instance, Rulke and Galaskiewicz (2000) found that groups composed of members with general knowledge outperform groups of specialists in centralized structures. Moreover, the performance of the members with general knowledge did not correlate with the group structure. In contrast, the performance of the specialists improved in decentralized groups. – We used agent-based modeling to highlight the advantages and disadvantages of several management styles in biology. Specifically, we compared the performance of centralized, hierarchical, organic, and egalitarian groups. In egalitarian groups all team members are connected with each other, while in centralized ones, they are only connected with the principal investigator. – We discovered that each group structure is associated with different epistemic challenges. While exact numbers are beyond the scope of agent-based models such as ours, results indicate the general trends when it comes to the impact of one-on-one meetings, the distribution of time spent on experimenting and communicating, and seminars on the performance of different group structures. – Our model shows that a weakly-connected organization of researcher in biology can improve the speed and reliability of the scientific pursuit when the epistemic space is complex. Furthermore, we found that an excessive amount of communication might encourage scientists to merely follow trends in the community, instead of pursuing their own ideas. Thus, egalitarian (fully-connected) teams perform best on simple epistemic spaces but underperform when it comes to exploring diverse hypotheses. These findings suggest that a supervisor who encourages diversity of thought, or a scientific community that promotes the exploration of alternative hypotheses, can improve the reliability of the scientific process. Moreover, the introduction of seminars to the model, changes the epistemic performance in favor of weaklyconnected teams (Figure 1). – All academic communication requires a relevant degree of certainty. Specifically, when it comes to one-on-one meetings an epistemic threshold for communication is beneficial, exemplifying the simple but effective rule of “thinking before speaking”. As for presenting, it seems that the evidence threshold has to be higher but not absolute (Figure 1). In this way, new ideas are shared early in the discovery process, enabling other researchers to build upon them. Finally, the reason why hierarchical groups might be considered as the best option from a pragmatic point of view, is that their overall performance under parameter changes remains solid, while other structures reach both peaks and low points.

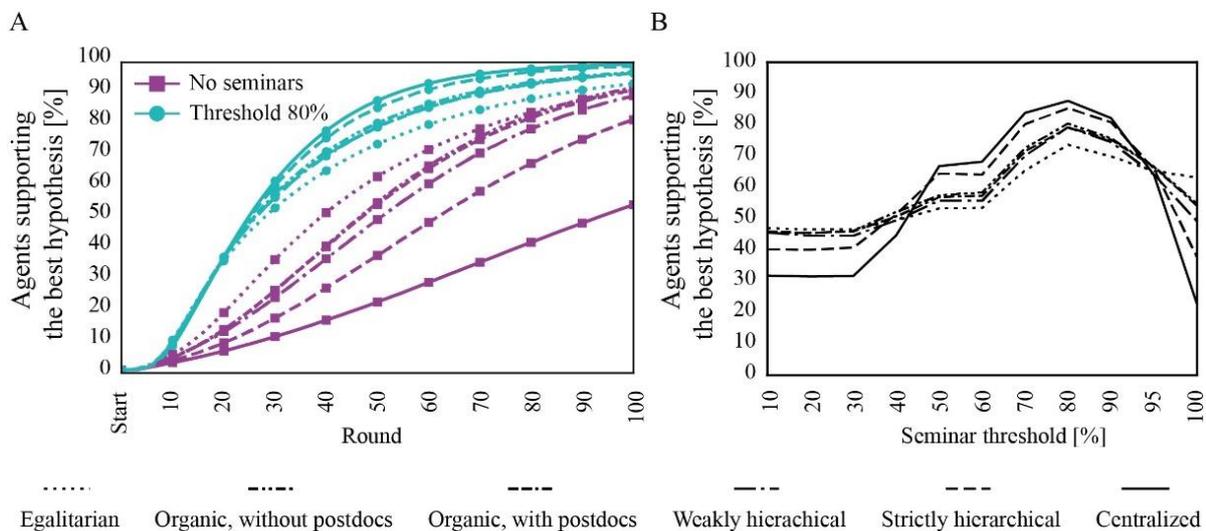


Figure 1: Finding the best hypothesis based on the seminar threshold. a) Standard model without seminars (purple) and model with seminars (blue). b) Agents supporting the best hypotheses after 50 rounds depending on the seminar threshold. When the agents can present their evidence to the whole network, the centralized and strictly hierarchical groups perform best.

Spagnesi, Lorenzo: “Idealization and Knowledge of Nature: A Kantian Approach”

All scientific models, including those used in physics, chemistry, biology, economics, and geology, contain idealizations, or assumptions that are known not to be true (Cartwright 1980; Nowak 1980; Giere 1988; Suárez 2009; Potochnik 2017). Scientists for example, assume frictionless planes, infinitely sized populations, and perfectly rational agents. Ideal assumptions such as these intentionally misrepresent the empirical system that is being studied, and yet they allow for the formulation of theories and laws of nature in a variety of scientific domains. While the benefits of idealizations are undeniable, philosophers of science have reached no consensus on how we can use ideal assumptions to know how things work in the real world (Weisberg 2007, 2012; Godfrey-Smith 2009). Some have argued that the inaccuracy of idealizations implies the falsity of scientific laws (Cartwright 1980, 1989; Lange 1993). For others, these are heuristic tools that may disappear with the advancement of science (Nowak 1980; Wimsatt 1987). A third position is that they are ineliminable, yet felicitous falsehoods that advance our understanding in a way similar to other forms of representations in art and literature (Frigg 2010; Elgin 2017). – During the Enlightenment, the philosopher Immanuel Kant developed a normative account of scientific investigation that can inspire a new approach to the contemporary debate. Kant argued that scientific investigation is possible only if guided by ideal assumptions—what he calls regulative ideas. These ideas are not true of any object of nature, and yet they are not mere heuristic tools or fictional representations of reality. They are necessary rules governing the construction and assessment of scientific explanations. To use some of Kant’s examples, the ideal assumption of fundamental power is the rule that allows to seek empirical explanations of various powers in psychology; the ideal assumption of organized beings is the rule that allows to seek physical explanations in biology (Kant 1998, 2000). In short, regulative ideas set the explanatory goals of scientific investigation. As such, they regulate the investigation of what reality is like and significantly contribute to our knowledge of nature. – In this paper, I suggest that Kant’s account of regulative ideas can help us reconcile the ubiquity of idealizations in contemporary science with a realist commitment to scientific knowledge. First, I develop a Kantian stance on idealizations that emphasizes the necessary, yet merely normative function of scientific idealizations. In other words, I argue that idealizations are not defined by their descriptive relation to

objects, but rather by the goals and standards they set up for empirical investigation. As such, idealizations are to be thought as explanatory ideals that are indispensable to the scientific investigation of nature. In the second part of the paper, I evaluate the benefits of a Kantian stance on scientific idealization on the basis of its capacity to: (i.) find an alternative to heuristic and fictionalist readings of scientific models; (ii.) explain the revisability and progress of scientific inquiry; (iii.) elucidate the shareability of scientific models by an epistemic community.

Stegenga, Jacob & Menon, Tarun: “A New Defence of the Value-Free Ideal”

The value-free ideal holds that scientific reasoning should not be influenced by non-epistemic values (social, political, or ethical values). It is on its face extremely plausible—like the ideal of world peace, one might think that the value-free ideal hardly needs philosophical defence. Science is our best guide to discovering objective facts about the world, and the value-free ideal is meant to block the intrusion of values which can bias the scientific process. Despite its *prima facie* appeal, the philosophical literature has articulated challenges to the value-free ideal, which has led to an emerging consensus that the value-free ideal is not attainable. Moreover, its status as an ideal is contested, with some arguing that value-freedom is not desirable for science. The ambition of this paper is to articulate a novel version of the value-free ideal which avoids the existing philosophical challenges. – The feasibility and desirability of an end can be decoupled from the feasibility and desirability of means for that end. Our core argument is based around the distinction between the end-state feasibility of an ideal and the pursuit feasibility of an ideal, and between the end-state desirability of an ideal and the pursuit desirability of an ideal. In general, a particular end may be unfeasible or undesirable, yet taking means to pursue that end may be feasible or desirable. The ideal of world peace may be impossible to achieve, yet pursuit of that ideal is both possible and good; de-escalating conflicts, disarmament, and a more equitable distribution of resources are all potential means to achieve world peace, and these means are, at least to some degree, feasible to enact. Same with desirability. The goal of working eighty hours per week is undesirable to us, but since we are lazy and work a mere ten hours per week, pursuing that end would be both possible and good. The existing challenges to the value-free ideal have focused on end-state feasibility and desirability rather than on pursuit feasibility and desirability. Yet, one can grant that the value-free ideal is neither end-state feasible nor end-state desirable, while maintaining that the value-free ideal is pursuit feasible and pursuit desirable. That is our goal. The conclusion of our argument is a specific—and as far as we know, novel—version of the value-free ideal, which holds that scientists ought to act as if science should be value-free. – We start by articulating the primary arguments in favour of the value-free ideal and the main challenges to the ideal (§2). We then argue that the value-free ideal is pursuit feasible (§3). In (§4) we argue that the pursuit desirability of an end can be decoupled from the desirability of the end itself, and we argue that the value-free ideal is pursuit desirable. We close with a modest criticism of the main argument against the end-state desirability of the value-free ideal, and then situate our defence of this new version of the value-free ideal among compelling views about values in science and the structure of scientific research (§5).

Sterkenburg, Tom: “The No-Free-Lunch Theorems of Supervised Learning” (Symposium “Learning from Data: The Secret to Success”)

The no-free-lunch (NFL) theorems of supervised learning (Wolpert, 1996; Schaffer, 1994) are an influential collection of impossibility results in machine learning. Computer scientists have ranked these results among the most important theorems in statistical learning (von Luxburg and Schölkopf, 2011), while some philosophers have read them as a radicalized version of Hume’s induction skepticism (Schurz, 2017). – In a nutshell, the results say—or rather, are usually interpreted as saying—that we cannot formally justify our machine learning algorithms. That is, we cannot formally ground our conviction that some learning algorithms are more sensible than others: that we have reason to think

some algorithms perform better in attaining the epistemic goals that we designed them to attain. In Wolpert's original interpretation, "all learning algorithms are equivalent," so that, for instance, a standard learning method like cross-validation has as much justification as anti-cross-validation (Wolpert, 2021). – Yet for many such standard learning algorithms we do seem to have a justification. The field of machine learning theory is concerned with deriving mathematical learning guarantees, that purport to show that standard procedures, like minimizing empirical error on the training set, are better than other possible procedures, like maximizing empirical error (Shalev-Shwartz and Ben-David, 2014). This raises a puzzle. How can there exist a learning theory at all, if the lesson of the NFL theorems is that learning algorithms can have no formal justification? – In this talk, I investigate the implications of the NFL results for the justification of machine learning algorithms. The main tool in my analysis is a distinction between a conception of learning algorithms as purely data-driven or data-only, as instantiating functions that only take data, and a conception of learning algorithms as model-dependent, as instantiating functions that, aside from input data, also ask for an input model. I argue that the NFL theorems rely on the former, data-driven conception of learning algorithms; and that there is here an important parallel to the philosophy of induction. – Namely, discussions surrounding the NFL theorems share a questionable presupposition with Hume's original argument for inductive skepticism: the idea that the performance of our inductive methods must be grounded in a general postulate of the induction-friendliness of the world. Contemporary philosophical work denies the cogency of such a principle, and advances a local view of induction (Okasha, 2005). This leads me to a local view of learning algorithms: the model-dependent perspective. Many standard learning methods, including empirical error minimization and cross-validation, take two inputs: data, and an explicitly formulated model or hypothesis class, which constitutes a choice of bias. What we can reasonably demand from such model-dependent algorithms is that they perform as well as possible relative to any chosen model. Consequently, learning-theoretic guarantees are relative to the instantiated models the algorithm can take, and it is in this form that there is justification for standard learning algorithms.

Šustar, Predrag – see Balorda, Vito

Šustar, Predrag – see Blečić, Martina

T

te Vrugt, Michael & Tóth, Gyula I. & Wittkowski, Raphael: "Irreversibility in statistical mechanics: from quantum mechanics to soft matter theory"

Finding an explanation for thermodynamic irreversibility is among the central problems in philosophy of physics. The microscopic laws governing the motion of individual particles are invariant under time-reversal, i.e., they do not allow to distinguish between past and future. Macroscopic thermodynamics, on the other hand, has a clear arrow of time characterized by the monotonous increase of entropy. Thus, there appears to be an inconsistency between microscopic and macroscopic physics. If a system is prepared in a certain "abnormal" initial condition, anti-thermodynamic behavior should be possible in principle. – In his book "Time and Chance", David Albert has suggested that this problem might be solved using the Ghirardi-Rimini-Weber (GRW) theory. GRW theory is a modification of quantum mechanics in which the wavefunction is assumed to undergo random collapses at a certain frequency. The objective stochasticity of GRW theory, Albert argues, allows to link the probabilities arising in

statistical mechanics to quantum-mechanical probabilities. In particular, the random collapses are supposed to restore thermodynamic behavior on the macroscopic level even if the system starts in an “abnormal” initial condition that would lead to anti-thermodynamic behavior in the classical case. – In this contribution, we present a computer experiment that is used to test Albert’s proposal. First, we prepare a many-particle system in an abnormal initial condition that leads to anti-thermodynamic behavior (the temperature difference between two systems in thermal contact spontaneously increases) if the system is modeled using classical mechanics. Then, we simulate the behavior of the system using GRW theory. Interestingly, we observe that the system still exhibits anti-thermodynamic behavior, i.e., the GRW collapses do not produce thermodynamic irreversibility if it is not already present in the classical case. The problem remains if we consider not standard GRW theory but its modern variants in which the collapse of the wavefunction is associated with a certain friction. Consequently, our simulations show that the GRW theory cannot explain thermodynamic irreversibility and thereby refute Albert’s proposal. – To make progress in the understanding of irreversibility, we take a closer look at the way irreversible dynamics is studied in modern physics. For this purpose, “the problem” of irreversibility is structured into five sub-problems (concerned with irreversibility in thermodynamics, the definition of equilibrium, coarse-graining, the approach to equilibrium, and the arrow of time). As an example, we consider dynamical density functional theory (DDFT), which provides a highly successful description of soft condensed matter systems. It is found that DDFT indeed provides valuable insights for the philosophical foundations of physics, since different “schools” in the philosophy of statistical mechanics correspond to different forms of DDFT. For example, deterministic DDFT belongs to the Gibbsian approach to statistical mechanics, while stochastic DDFT belongs to the Boltzmannian approach. Moreover, DDFT for colloidal fluids explains irreversibility in an interventionistic framework, whereas DDFT for atomic fluids requires coarse-graining. Consequently, DDFT shows that these different philosophical views are not mutually exclusive, but are all necessary to fully understand the modern practice of statistical mechanics.

Toader, Iulian: “Einstein Completeness as Categoricity”

Einstein's argument for the incompleteness of quantum mechanics, which did not make it into the EPR paper (Einstein, Podolski, and Rosen 1935), but was clearly formulated in his letters to Schrödinger and several subsequent papers (Fine 1986), has been suggested to deploy a semantic notion of completeness as categoricity (Howard 1990). The suggestion is motivated by the view, defended by Einstein, that quantum mechanics fails to assign a unique wavefunction to the same physical state of a system, since the assignment depends on the measurement performed on a spacelike separated but previously interacting system. If there are distinct wavefunctions for the same system, and if one is justified in considering them as non-isomorphic models in the sense of formal semantics, then this is enough to show that quantum mechanics is incomplete, i.e., non-categorical. – The present paper offers a rigorous reconstruction of this line of thought, thereby justifying an understanding of Einstein completeness as categoricity. The key conditional claim to be carefully articulated and defended is the following: “if one understands a theoretical state as, in effect, a model for a set of equations plus boundary conditions ..., then Einstein's conception of a completeness requirement should really be understood as a categoricity requirement.” (Howard 1992, 208) To do this, I start from the notion of an EPR state, formally defined for a composite quantum system (Arens and Varadarajan 2000), but suitably adjusted to an algebraic setting, and I argue that Einstein completeness fails due to the unitary inequivalence of non-regular Hilbert space representations of the Weyl algebra describing one of its subsystems. The argument, as I will discuss in some detail, assumes unitary equivalence as a necessary (though not sufficient) condition for categoricity (Toader 2021). – Incidentally, my reconstruction shows not only that Einstein's argument is far from being “muddled” (Landsman 2006, 234), but also that it is far from being as dull as one sometimes takes it to be, when one claims that it entails “overcompleteness” (Lehner 2014, 319), in the sense of theoretical underdetermination, rather than the incompleteness of quantum mechanics. – Furthermore, I think that reading Einstein completeness as categoricity throws new light on the Bohr-

Einstein controversy. Bohr's complementarity doctrine has been interpreted in terms of the unitary inequivalence of non-regular Hilbert space representations (Halvorson 2004, but see Feintzeig et al. 2018) which formally vindicates a common view that Bohr's notion of completeness was significantly distinct from the descriptive completeness articulated and challenged in the EPR paper (Norsen 2017, 148). Indeed, it turns out that the semantic sense in which Bohr thought quantum mechanics was complete is precisely the sense in which Einstein argued it wasn't. More exactly, whereas Einstein demanded categoricity, Bohr maintained that any attempt to satisfy this requirement would go against complementarity.

Tóth, Gyula I. – see te Vrugt, Michael

Tramacere, Antonella: “Has the evolutionary study of the mind reached an impasse?”

Classical theories of cognition distinguish general and specific information-processing, respectively called domain-general and domain-specific systems. An explanation based on domain-specificity (i) isolates some function or mechanism within the system of interest, and (ii) delineates the functional domain of that function or mechanism – i.e., its inputs, internal processes, and outputs. On the contrary, a domain-general system applies a general type of function to potentially any kind of input, and works in potentially any venue. – The distinction between domain-specific and domain-general has been borrowed by early evolutionary psychology, pervading discussions of the evolution of the human mind. Domain-specific mechanisms became the result of evolutionary adaptations, producing constraints in cognitive processes all the way up, from genes to the mind. On the other hand, domain-general systems are seen as learned, plastic and context-dependent mechanisms, which result from socio-cultural dynamics during development. – However, much of the recent literature has weakened this polarization in one way or another (see for example Heyes, 2018; Barrett 2005). Domain-general processes of learning can also produce domain-specific mechanisms, as in the case of literacy. The cognitive mechanism of literacy is isolable and decomposable in specific sub-mechanisms; plus, it is neurally localized. However, literacy is not an adaptation by natural selection, as it emerged quite recently in human evolution, and it is learned culturally. In line with Heyes (2018), call cognitive gadgets those mechanisms that result from the evolution of cultural learning, rather than genetic evolution. – Research has also shown that mental mechanisms can be acquired through domain-general learning during ontogeny, and yet be the result of a process of natural selection on genes. Genetic evolution does not necessarily produce genetically programmed, cognitive specialised neural structures, but can also enhance plasticity and general learning possibilities. This is the case of vocal imitation, which evolved genetically (through the mutation of the Foxp gene family), while increasing the general-learning capacities of the species. Call cognitive twists mechanisms evolved through genetic evolution, but which are still flexible and culturally learned. – If cognitive gadgets and cognitive twists exist, then the individuation of domain-specific systems can hardly say anything on the evolutionary origin of mental mechanisms. If domain-specific systems can be both adaptations and the result of domain-general learning, then functional analysis of neural mechanisms cannot provide conclusive information on the phylogenetic history of those systems. Further, if the criteria proposed by evolutionary psychology to identify mental adaptations are inadequate, then the evolutionary study of the mind has reached an impasse. – I provide a set of criteria for overcoming this impasse, including a) Theoretical Model identifying a plausible selection scenario in the biological and/or cultural domain explaining the evolution of the mechanism; b) empirical studies identifying distinct neurobiological mechanisms co-varying functionally with the mental mechanism (Evidence of Function) and c) evidence of evolutionary changes (both at the genetic and the paleoanthropological level) which are compatible with natural selection scenarios (Evolutionary Change). – I contend that these criteria do not oversimplify the

complexity of the interactions between learning and evolution, and scatter the polarization which looms over the debate on the evolution of human cognition.

Trappes, Rose: “The Pervasiveness of Sex in Behavioural Ecology”

Sexual difference is pervasive in behavioural ecology. It appears in theories about mate choice and sexual signalling, it is factored into most data analyses, and it is itself a topic of empirical research. There have been important feminist critiques of sexual selection theory, especially in primatology (Hrdy 1986; Haraway 1989) and sociobiology and evolutionary psychology (Lloyd 1993; 2005; Martin 2003), and there are of course decades of discussions about defining sex in humans. Yet there has been little investigation into the concept and theoretical role of sex in behavioural ecology more broadly, beyond humans and their near primate relatives. In this talk I consider how sex is defined in behavioural ecology, what roles it plays, and why it is so pervasive. – As evident in practices of sexing (determining the sex of animals) and theoretical modelling, behavioural ecologists define sex based on morphology, reproductive roles, and behaviour. I suggest a general definition of sex in behavioural ecology, according to which reproductive role is primary—females produce eggs, males sperm. – Though it is sometimes a topic of research in its own right, sex most frequently appears in behavioural ecology as a way to explain other phenomena. Sex is often factored into data analysis as a way to explain variation, such as the variation within a group in resource use rates or social interactions. In addition, sexual selection and sexual signalling theories take sex as important for explaining broader evolutionary changes in morphological and behavioural features, population structure, and so on. I therefore argue that sex’s main role in behavioural ecology is to serve explanatory goals. – Finally, I will make some suggestions about why sex plays this explanatory role so often. First, in many species sex is a difference that is discrete and relatively easy to identify. This makes it a useful way to explain variance, in contrast to more continuous or less obvious differences that require careful measurement and statistical procedures. Interestingly, however, behavioural ecologists apparently still try to use sex to explain variation even when sex is hard to determine, as in species that are not sexually dimorphic. Hence, there must be other reasons for the pervasiveness of sex.

Trotter, Frida: “Breaking underdetermination with norms”

In this paper I defend the view that some cases of underdetermination of theories by evidence are “broken” via resort to arguments of normative nature. This may stand in opposition with the idea that theory choice would always be based exclusively on “objective” grounds, for instance via an assessment of a theory’s superempirical virtues (Schindler, 2018). Interestingly, holding different normative assumptions with respect to what science ought to do may motivate different views of science altogether. In order to expose my argument in detail, I focus on the underdetermination of different interpretations of quantum mechanics, and refer to examples of normative assumptions defended by some of their prominent proponents. – For a leaner but not less general argument I consider two realist interpretations of QM, namely, the Many Worlds Interpretation (MWI) (Wallace, 2012) and Bohmian Mechanics (BM) (Bohm, 1952 and Holland, 1993). Both hold a realist view of the theory as referring to a particular ontology, and both provide a clear answer to the request for explanation of the quantum phenomena. These two interpretations are underdetermined by evidence, as no experiment is currently able to decide for, or to rule out, either ontology, but they are *differently virtuous*. The MWI is formally simple, and the solution it provides to the “preferred basis problem” — the problem of why we seem to experience only facts belonging to a single branch of the universe — has been criticised for being ad hoc. BM on the other hand is ontologically simple, and it includes ad hoc postulates such as the equilibrium hypothesis and the notion of quantum potential. Both theories are empirically adequate, internally consistent and fruitful, as they give rise to new predictions — these latter however have been either

unsatisfactorily tested or not tested at all. – Given this different array of virtues, one may wonder whether we could in fact assess which theory is overall more virtuous, and logically select it as superior to the other. The problem is precisely that the matter is not so simple: one can in fact *disagree* with respect to which virtue is more desirable or more important than another. And it is here that norms are called into account. Despite it may not be possible to grade theories merely on the basis of their virtues, the supporters respectively of the MWI and of BM do argue for the superiority of their interpretation over the others. In this paper I provide examples of disagreement in these two camps with respect to what science ought to do, and of how these different claims underlie the preference for a particular interpretation of QM over the competitors. I also argue that by taking further these different commitments about what science *ought to do*, different portraits of what science *is* may start looming in the background.

Trotter, Frida – see Frembs, Markus

V

Valkovic, Martin: “Cultural evolution of human cooperation” (Symposium “The Legitimacy of Generalizing Darwinism”)

Cooperation is central to our species’ way of life, in all its domains and on a uniquely large scale. We cooperate to obtain food, raise children, fight forest fires, prevent (and enable) nuclear holocaust and organise raves. Unlike in other species, human cooperation is not limited to close kin nor is as temporally and spatially limited. It is also by no means a recent development. Biological approaches to the evolution of cooperation abound, but they are only a part of the story: there is a broad consensus on the importance of cultural inheritance and selection in explaining the change in our species over time. – I will present two groups of theories of cultural evolution, termed cultural evolution1 and cultural evolution2. Cultural evolution1 accounts analyse the evolution of cooperation in terms of transmission and selection of traits, norms, behaviours, social structures and similar in a predominantly non-genetic manner. Here, the entities making up an evolving population are biological individuals or groups of them, and culture is a part of their phenotype (Godfrey-Smith 2009). Cultural evolution2 accounts, on the other hand, take cultural phenomena themselves to be undergoing evolution, as a consequence of varying and reproducing differentially. That is, cultural variants, such as social organisations, institutions, norms, social roles, identities, and the like, or groups of them constitute a Darwinian population here (Godfrey-Smith 2009). – As will become apparent, these two groups of theories differ considerably both between each other and within themselves, and I will argue there is a need to compare them and to determine their respective places in the story of the evolution of human cooperation. On the other hand, theories from both groups need to be compared to Darwinian biological evolution theories, to determine which elements are shared, and if and where the theories go apart.

van Panhuys, Marianne & Hillerbrand, Rafaela: “Epistemic risks and computer simulation: a case study from particle physics”

In philosophy of sciences, the issue of epistemic risk is usually addressed in terms of inductive risk, focusing on the process of decision-making to accept or reject hypotheses based on empirical evidence. This topic is widely discussed in the literature on the Argument from Inductive Risk (AIR) (Steel, 2010) and mainly concerns with the role of value-laden judgements in weighing evidence to prevent from social and ethical harm. – In many sciences today, however, empirical reasoning is highly inferential as experiments rely on complex instrumented disposals. This means that there is a long process before confronting evidence to hypothesis. This process often involves an increasing use of computer simulations, may it be in life science or particle physics where computer simulations are, for example, centrally involved in the design of particle detectors and data generation. The crucial role of these computer-based practices, which are in this context precondition for empirical reasoning, call for further philosophical insight regarding risks. – In this paper, we zoom in on particle physics and aim to expand the framework of epistemic risks to particularly address the issue of computer simulation-related risks. Based on a case study from ATLAS experiment in top-quark physics we argue that there are relevant epistemic risks besides inductive ones that go beyond social and ethical impacts. The subsumption of risks under inductive ones is insufficient to address the variety of risk arising in the course of scientific inquiry as well as to address the collaborative feature of producing scientific knowledge (Biddle & Kukla, 2017). After analyzing contingent choices made in the experimental process, we propose to frame epistemic risk as the risk to not fulfil one’s epistemic aim, distinguishing between local (e.g., prediction) and global (e.g., discovery) aims. Our contribution can be understood as an attempt to locate uncertainty and risks and explicate relationships at stake.

Vogt, Lisa – see Hirèche, Salim

W

Wachter, Tina: “Does Referencing in QM Require Free Logic?”

In the philosophical debate concerning the individuality and distinguishability of quantum particles, there are three main defence strategies of the Principle of the Identity of Indiscernibles. I will focus on the summing defence strategy regarding fermionic non-entangled and entangled states. – According to the summing defence strategy, in a fermionic GMW-entangled state, there are no numerically distinct entities, but only a unified, undivided whole that can be (descriptively) referred to. As I will show, descriptive referencing holds for the non-entangled as well as the entangled situation, whereby free logic helps us to understand how we can properly refer to particles (i.e. similar fermions) before unification, while entangled, and after dividing the physical system again via an EPR-like measurement.

Weber, Marcel: “Modeling Modality: The Case of Evolvability in Evo-Devo”

Biological modalities, i.e., biologically possible, impossible or necessary states of affairs, have not received much attention from philosophers. Yet, it is widely agreed that there are biological constraints on physically possible states of affairs, such that not everything that is physically possible is also biologically possible (while, of course, everything that is biologically possible is also physical possible). For example, while a flying elephant is both physically and biologically impossible, an elephant with feathers is physically possible but biologically impossible. Another widely agreed feature of biological

modalities is their relativity to a specific time and lineage, and their gradable character. Other than these basic feature, biological modalities remain ill-defined and there are only scant attempts to give an explicit account of them (Dennett, 1995; Hindermann, 2016). – In this talk, I investigate what kind of modality underlies the concept of evolvability in evolutionary developmental biology or “evo-devo”. This concept tries to capture the capacity of an organism or a lineage to sustain genetic changes that enable it to evolve or to evolve adaptively. Thus, evolvability is to evo-devo what fitness is to natural selection theory. Philosophers of biology have construed both evolvability and fitness in dispositional terms, or as a probabilistic propensity. As there are known difficulties with these notions, it is tempting to try to give a direct modal account of evolvability, i.e., an account that does not rely on either dispositions or propensities. – The basic idea of the proposed approach is to construe evolvability as a kind of accessibility in a modal space. Obviously, the difficult part is to specify this modal space. While there may not be a general way of defining such a modal space, there exist model systems for which it may be possible. One such model system consists of small RNA molecules that can fold into distinct secondary structures which are determined by intra-molecular base pairing. It has been shown that what matters for such RNAs being able to evolve a different secondary structure is the presence of so-called neutral networks, that is, sets of sequence modifications that do not affect the secondary structure (Fontana, 2002). Molecules can walk such networks by random mutations and eventually switch over to a neighboring network. – What makes it possible in this example to model the modal space around a given RNA sequence is the availability of a detailed genotype-phenotype map or GP-map that shows what phenotypes can be obtained from what genotypes. This map can be used to define a non-metric neighboring relation that determines accessibility. Thus, the modal space in such cases is quite distinct from those constructed by philosophers, e.g., David Lewis’s similarity metric for possible worlds.

Wenz, Daniel – see Kasprovicz, Dawid

Wilholt, Torsten: “Symmetries and Asymmetries in Epistemic Risk Management”

It is often assumed that scientists may (and should) consider the extra-scientific costs of getting it wrong in their methodological decisions. In contradistinction, it is argued, the extra-scientific benefits of getting it right may and should not be thus taken into account. (Heather Douglas argues so explicitly, and the view seems to be shared by other authors.) I will argue in favor of treating the costs of getting it wrong and the benefits of getting it right symmetrically: In cases of decisions in which scientists are obliged (or permitted) to consider the costs of getting it wrong, the situation is such that scientists are equally obliged (respectively permitted) to consider the benefits of getting it right. But even if an asymmetry between the two cannot be justified, the search for its roots can lead us to other asymmetries that may be worth discussing in the context of the ethics of epistemic risk management. I will argue that in many stock examples, the apparent asymmetry between taking the seriousness of mistakes into account and considering the benefits of getting it right is related to the moral asymmetry between action and omission. An inquirer who has not adequately taken into account the possible harms of getting it wrong and who then actually gets it wrong and thereby induces these harms can be understood to have failed his moral responsibilities by actively causing harmful consequences. On the other hand, an inquirer who has not properly taken the benefits of getting it right into account and as a consequence does not establish the result that would have brought about good effects seems “only” to be culpable for an omission. The intuitive plausibility of a morally relevant difference between doing harm and failing to prevent it is quite strong, as non-consequentialist ethicists have often pointed out. It explains much of the reluctance against putting the consequences of error and the consequences of true results on a par when it comes to epistemic risk management. I will then discuss whether grounding the relevant asymmetry in the

doing allowing distinction actually withstands philosophical scrutiny. I will explore both Jonathan Bennett's distinction between behaviors that are positively relevant to an outcome from such that are negatively relevant and Philippa Foot's distinction between rights to non-interference and rights to goods and services as resources for understanding morally relevant differences between doing and allowing. The conclusion of my discussion will be that for the case of epistemic risk management, no convincing morally relevant difference of this kind exist. - There may be rare, extreme cases where the production of a certain kind of error itself constitutes the infringement of someone's right to non-interference. But even in these cases, the advisable action would be to abstain from conducting the inquiry completely. A different treatment of the consequences of error and consequences of true results within epistemic risk management is never morally called for.

Williamson, Jon: "Applying Evidential Pluralism to the Social Sciences" (Symposium "Mechanisms in the Cognitive and Social Sciences")

Since around the year 2000, philosophers of science have produced a great deal of interesting research on the role of mechanisms in science. One strand of this research concerns the role of mechanistic evidence in establishing causal claims. Russo and Williamson (2007) argued that in the biomedical sciences, a causal claim is established by establishing (i) that the putative cause and effect are correlated, and (ii) that there exists a mechanism linking the two which can account for this correlation. This thesis has the following important consequence: while quantitative studies (in particular, randomised controlled studies) provide excellent evidence of correlation and, in the right circumstances, can provide evidence of the existence of a mechanism, it is important to also consider other evidence of mechanisms when assessing a causal claim. This motivates a kind of Evidential Pluralism. In medicine, this form of Evidential Pluralism has led to a proposed modification to evidence-based medicine, called EBM+. Parkkinen et al. (2018), for instance, developed procedures for evaluating mechanistic studies alongside clinical and epidemiological studies, when assessing the effectiveness of an intervention or when ascertaining the effects of exposure to an agent. – This paper argues that Evidential Pluralism applies equally to the social sciences. In the social sciences, as in the biomedical sciences, establishing causation requires establishing both correlation and mechanism---social mechanisms, in this case. While quantitative association studies can provide some evidence of mechanisms, in addition to good evidence of correlation, other sorts of study also provide good evidence of social mechanisms--notably, certain qualitative studies. – We argue that there is scope to apply Evidential Pluralism to the social sciences. First we show that the lessons from evidence-based medicine can be carried over to evidence-based policy, and that Evidential Pluralism can provide an account of the assessment of evidence in evidence-based policy. We compare this account to that provided by realist evaluation, which also has a central role for mechanisms. Second, we use case studies to argue that Evidential Pluralism additionally applies to more theoretical social sciences research, and can be used to elucidate the confirmation relations in basic social sciences research. Third, we show that Evidential Pluralism can provide new foundations for mixed methods research, because it offers a precise account of the need for mixed methods when establishing causation in the social sciences. – We then respond to two objections to the claim that Evidential Pluralism can be applied to the social sciences: one due to Julian Reiss and a second due to Francois Claveau. We conclude that Evidential Pluralism has much wider scope than originally envisaged, and sheds new light on the use of evidence in the social sciences.

Wittkowski, Raphael – see te Vrugt, Michael

Wray, K. Brad: "The Epistemic Significance of the Size of Research Teams"

Boyer-Kassem and Imbert published a paper with the provocative title: “Scientific Collaboration: Do Two Heads Need to be More than Twice Better than One?” Their paper invites us to think about whether increasingly larger research teams are superior to small teams. I believe that this is the next frontier in the social epistemology of collaborative research. We need to develop a better understanding of how research teams behave in order to understand how they can produce knowledge. We want to understand what qualities of research teams — for example, (i) their size, (ii) their physical distribution across institutions, (iii) their constitution in terms of members, or (iv) their persistence — contribute to or impede our effective pursuit of knowledge. – In this paper, I make a first step in advancing our understanding of the epistemic significance of the size of research teams. I begin by reviewing some recent empirical research that provides insight into the issue. Then I present some new data that provides evidence of differences in performance of research teams of different sizes. The data draw attention to the relative propensity of different sized teams to produce research publications that need to be retracted. For example, though smaller research teams, teams of two to four scientists, retract papers at a proportionally higher rate than larger teams, that is, teams of 12 or more scientists, when larger teams do retract papers, they often do not cite the reason for the retraction. In fact, larger teams are more inclined to publish ambiguous retraction notices. Such retraction notices give us reason to believe that they are hiding misconduct, or are incapable of identifying the causes of the problem that led to the retraction. – Finally, I draw some normative conclusions about team size, relating my findings to other recent research on team size. Given the data I present, there are grounds for believing that there is no optimal size for a research team. That is, research teams of all sizes face challenges, and no particular team size ranks higher with respect to all the various measures. This is not, though, as different research problems are likely to require different sized teams. – Further, given the differences in the behaviors of research teams of different sizes, in particular, given the ways in which they differ with respect to how they deal with the need to retract a published article, I argue that we have reason to believe that the results of research produced by larger teams are less reliable in some sense. Because we are less certain about the ability of such teams to resolve issues adequately when problems arise, we have less reason to trust such teams in general. The data thus provide an important corrective to the increasing trend to conduct scientific research in larger teams.

Y

Yu, Li-An: “Epistemic injustice of climate change: the coherence problem of specific and general information”

In this talk, I present an epistemic challenge about climate change, and then a potential remedy.

1. The coherence problem of specific and general information. Studies about the climate produced by scientists, in particular those of the IPCC, have been considered authoritative and invoked for political decision-making in a global context (Shepherd & Sobel, 2020). However, this assumption has created an epistemic situation, where some specific populations can challenge global claims about the climate on their own epistemic grounds. I term it the coherence problem of specific and general information. For instance, it is known that the current climate projections of GCMs have a higher degree of uncertainty regarding the dynamics of precipitation patterns in the tropics, and are unable to discern the distribution of precipitation within small regions such as London and Taiwan. Such limited resolution could be of great practical significance, as people living in these areas might find that global claims about local precipitation patterns do not agree with their on-site experience. – Moreover, when one assumes claims that the differences in carbon emissions are relevant to the 1.5-2.0 deg global mean temperature and the rise of the sea-level, Uyghurs and Himalayans might find it difficult to understand how the climate crisis can occur. The reason is that their everyday experience includes daily temperature oscillations of more than 20 deg C, and their settlements are thousands of kilometers away from the sea shore. Some indigenous groups in the Americas might feel ridiculed as the future climate scenarios are

described by experts to be catastrophic, while their cultures have already perceived the present world to be dystopic, as compared to their world before European settlement (Whyte, 2018). – Therefore, such contrasts between knowledge produced by specific experiences, geographies, histories, etc., and information derived from global models exemplify the coherence problem described. Local populations realizing this incoherence may come to believe that a global target of reducing carbon emissions alone has little connection to the problems they are facing, which may impede action.

2. Epistemic injustice about the climate. This situation exhibits epistemic injustice, meaning that the knowledge and experience of some competent and insightful persons are not considered relevant because of their social status, gender, race, etc. (Fricker, 2007). Climate modeling, for instance, is a practice exemplifying epistemic injustice of a kind, as the construction of models has from its beginning favored narrow modeling expertise. In fact, the adequacy and relevance of these studies may be cast into doubt by considering the specificity of experiences, geographies, histories, etc. Seeking cooperation with indigenous populations for climate action may be offensive as Westerners tell them how to save themselves from the climate crisis as if they were victims to be saved. Their hermeneutic frameworks seem to be thus disregarded.

3. An HPS approach to science communication. In this talk, I propose and emphasize a HPS approach to science communication, and clarify its potential to address the coherence problem. The results are expected to remedy the injustice that hurts both scientists and the people.

Z

Zaffora Blando, Francesca: “Merging of Opinions for Computable Bayesian Agents and Algorithmic Randomness” (Symposium “Learning from Data: The Secret to Success”)

A standard objection to subjective Bayesianism is that its reliance on subjective priors threatens the objectivity of scientific inquiry. A standard Bayesian response to this charge is that there is a sense in which prior probabilities are immaterial for the purpose of successful inductive learning: provided that certain conditions are met, the influence that diverging subjective priors have on posterior probabilities is guaranteed to be eventually washed out by the cumulating evidence. This response crucially relies on a family of results in the foundations of probability and statistics known as merging-of-opinions theorems. In a nutshell, merging-of-opinions theorems establish that, as long as their respective priors are sufficiently compatible to begin with, two (or more) Bayesian agents with differing initial degrees of beliefs are guaranteed to almost surely reach a consensus with increasing evidence. Thus, objectivity can be recovered in the form of inter-subjective agreement. – It is important to stress that almost-sure inter-subjective agreement does not happen in all circumstances: as mentioned above, its attainment requires some amount of compatibility among the initial beliefs of the members of a given community. For instance, the most well-known merging-of-opinions result, the Blackwell-Dubins Theorem (Blackwell and Dubins, 1962), shows that Bayesian conditioning leads to a strong form of consensus whenever the agents agree on probability-one events: i.e., provided that their respective priors are mutually absolutely continuous. However, absolute continuity is a rather strong form of compatibility between credences. It is thus natural to wonder whether merging of opinions, and what type of merging of opinions, can be achieved with weaker assumptions. – In this talk, I will address this question from the perspective of computationally limited Bayesian agents: that is, agents whose priors are computable probability distributions. I will argue that, for computable Bayesian learners, it is natural to appeal to the theory of algorithmic randomness to define notions of compatibility between priors—where algorithmic randomness is a branch of computability theory aimed at characterising the concept of an effectively

typical outcome of a given generating probability distribution (see, for instance, (Nies, 2009) and (Downey and Hirschfeldt, 2010)). In particular, I will show that the proposed notions of compatibility defined in terms of agreement on algorithmic randomness correspond to restricted forms of absolute continuity. Then, I will investigate different types of merging of opinions—such as the weaker notion of merging proposed by Kalai and Lehrer (1993), which only requires reaching a consensus over finite-horizon events—and show that, in general, the forms of compatibility between priors induced by algorithmic randomness lead to the attainment of inter-subjective agreement with probability one. Ultimately, the goal of this work is to gain a deeper understanding of exactly how similar the initial credences of more realistic, less-than-ideal Bayesian agents have to be in order for their posterior credences to eventually align.