Mark Addis. Categorical Abstract Model Theory and the Syntax of Scientific Theories There is a problematic discrepancy between current ways of characterising the semantics of mathematical theories and much philosophical thinking about whether scientific theories are best conceived as in semantic or a syntactic way. In the syntactic approach theories are analysed as deductive axiomatic systems (using quantified first order logic plus various relational extensions) in conjunction with appropriate empirical interpretations of non-logical terms [Carnap 1939 and Nagel 1961]. This approach was criticised for either ignoring or distorting many aspects of theory construction in science [van Fraassen 1980] and increasingly fell out of favour as logical positivism waned. It was gradually replaced by the semantic approach to scientific theories which held that theories are abstract specifications of a class of models where a model is a structure in which a theory is true [Suppe 1989]. The syntactic approach is still widely considered problematic and although work by Halvorsen Tsementzis [forthcoming] and Lutz [2015] is gradually making it more acceptable. Much of the reason for favouring the semantic approach stemmed from regarding first order logic as the paradigmatic form of logic with the greatest explanatory clarity. It will be argued that the syntactic approach was and is unpopular for reasons which have much more to do with the limitations of first order logic than any fundamental philosophical deficiencies of the position itself. In particular philosophical analysis which regards first order logic as the paradigmatic form of logic characterises the relationship between logic and the structure of scientific theories as it was rather than currently is or could be.

Within logic and mathematics there has been a gradual but fundamental change in how the semantics of mathematical theories are approached due to a number of related developments. The study of logic beyond the confines of the first order realm has been going in earnest since at least the late 1950s (such as with Mostowski's work on generalised quantifiers) and resulted in the general growth of abstract model theory. This growth marked a decisive shift away from a focus upon the isolated investigation of specific especially first order logical systems to one upon the relationships between a wide range of logics. Early abstract model theory assigned ascribed a central place to the notion of an abstract logic (with major results being Lindström's theorem [Lindström 1969] and Barwise's axiomatisation [Barwise 1974]) but made limited use of ideas from category theory. In the 1980s computer scientists Goguen and Burstall introduced the theory of institutions in order to relate various logics, such as fragments of many sorted first order logic and higher order logic with polymorphic types [Goguen and Burstall 1984]. The concept of an institution is more general than that of an abstract logic in that it achieves independence from actual logic systems through a fully categorical abstraction of the main logical concepts of signature, sentence, model, and of the satisfaction relation between them [Diaconescu 2012].

In the definition of an institution a category C consists of:

- a category Sign, whose objects are called signatures (of the various vocabularies) and whose arrows are called signature morphisms
- a functor Sen: Sign→Set, giving to each signature a set whose elements are called sentences over that signature,
- a functor Mod: Sign→ Catop giving for each signature ∑ a category whose objects are called ∑-models, and whose arrows are called ∑-model morphisms and
- a relation $\models_\Sigma \subseteq |Mod(\Sigma)| \times Sen(\Sigma)$ for each $\Sigma \varepsilon |Sign|$ called ∑-satisfaction, such that for each signature morphism $\phi: \Sigma \rightarrow \Sigma'$ the satisfaction condition $m' \models_\Sigma Sen(\phi)(e)$ iff $Mod(\phi)(m') \models_\Sigma e$ holds for each $m' \varepsilon |Mod(\Sigma')|$ and each $e \varepsilon Sen(\Sigma)$

The theory of institutions is not concerned with the internal structure of particular objects (which are in this case particular logical systems whether first order or not) but instead with how objects are related to other objects by homomorphisms. The institutional approach is well suited for the definition and classification of algebraic and geometric theories and structures. The flexibility which makes it useful for characterising these kinds of theories and structure is also very helpful for permitting the study of scientific theories from both the syntactic and semantic perspectives. The theory of institutions addresses Hempel's [1970] concern about to how study scientific theories which simultaneously require the use of more than one logical vocabulary as it allows the use of more than one formalism at the same time.

Since the theory of institutions also supports automated theorem proving it can be employed for automating the process of generation and refinement of scientific theories and models. Given these considerations there are substantial grounds for a reassessment of the value of engaging in the study of syntactic features and for questioning whether the current widespread rejection of syntactic analysis of scientific theories is well justified.

Vincent Ardourel. Boltzmann-Grad limit and irreversibility The derivation of the Boltzmann equation (BE) from the Hamiltonian equations of motion of a hard spheres gas is a key topic on irreversibility in statistical mechanics. Recent discussions focus on Lanford's rigorous derivation of the BE which is "maybe the most important mathematical result of kinetic theory"(Villani 2010, 100). However, the origin of irreversibility in this derivation is still unclear. According to Uffink and Valente (2015, 432):

"We discussed the problem of the emergence of irreversibility in Lanford's theorem. We argued that all the different views on the issue presented in the literature miss the target, in that they fail to identify a time-asymmetric ingredient that, added to the Hamiltonian equations of motion, would obtain the Boltzmann equation."

In this paper, I claim that Uffink and Valente are wrong about one of the different views in the literature, which is the role of the Boltzmann-Grad (B-G) limit, i.e the limit of infinitely diluted gases. Although this limit is sometimes viewed as a main ingredient in the appearance of irreversibility (e.g Valente 2014, 319), Uffink and Valente recently mitigate this account and suggest that the B-G limiting procedure "is not sufficient for the appearance of irreversibility" (2015, 424). I argue that their argument, which is based on a comparison made by Lanford (1981, 75) between the derivations of the BE and the Vlasov equation (VE), is misleading. My point is that, contrary to what Uffink and Valente suggest, the derivation of the VE does not require the B-G limit. One thus cannot draw any conclusions about the B-G limit based on an analysis of the derivation of the VE. Instead, I still argue that the B-G limit is a main ingredient to explain how a time-reversal non invariant equation (i.e the BE) can be derived from a set of time-reversal invariant equations (i.e the Hamiltonian equations of motion of a hard spheres gas).

The VE describes the evolution of the density of distribution in phase-space that a particle is located at the position q with velocity v when the interaction between particles is given by a sum of two-body potentials of the form $\varphi(N)(q1 - q2) = 1/N\ \varphi0(q1 - q2)$. It is derived from the Hamiltonian equations of motion of N particles when N—> ∞. Although the derivation of the BE and the VE look very similar, there is no appearance of irreversibility in the derivation of the VE since the H-function, occurring in the H-theorem, is constant with time. This contrasts with the case of the BE for which the H-function monotonically decreases with time. Uffink and Valente thus suggest that the B-G limit is not sufficient for the appearance of irreversibility.

I show that the derivations of the BE and the VE are actually based on two different limiting regimes. The VE is derived in an effective field limit, which only requires N—> ∞ limit. Instead, the derivation of the BE requires in addition that a—> 0 where Na2 converges to a finite quantity, which is strictly the B-G limit. This difference in limiting regimes is crucial because it explains the difference between the BE and the VE with regard to irreversibility. In the case of the derivation of the BE, the dynamics of the N hard spheres is no longer deterministic when the diameter a of spheres goes to zero: after each collision, the direction of particles is random (Norton 2012, 218; Golse 2014, 36). It is no longer possible to determine past trajectories. However, the diameter a is not a limit parameter in the derivation of the VE. This equation is derived from a model of N mass points for which particles have already zero diameter.

One may object that the BE can also be derived form a model of N mass points (Gallagher et al. 2013). But in this case, I stress that this derivation still requires the B-G limit: It requires N—> ∞ limit and a—> 0 limit for which the parameter a now corresponds to the range of the potential. This contrasts with the derivation of the VE where there is not such limit parameter.

Finally, I make clear that the use of these two different limiting regimes is based on two different models about how particles interact with each other. The derivation of the BE assumes a strong but local coupling between particles, which corresponds to a collisional model. Instead, the derivation of the VE assumes a weak but global coupling between particles, which is a model where each particle interacts with other particles without colliding (Golse 2003, 2). These two limiting regimes are two ways to make compatible the use of an infinite limit for Hamiltonian systems with the constraint that the average energy per particle remains bounded, which can be done "either by scaling the strength of the force, or by scaling the range of potential"(Gallagher et al. 2013, 7). The first case leads to the VE and the second one, to the BE.

Bengt Autzen. Musing on Means The relationship between expected offspring numbers and future offspring numbers has been widely discussed in the philosophical literature on Darwinian fitness. This is due to the fact that a prominent view on the nature of fitness - the 'propensity interpretation of fitness' - initially adopted the expected number of offspring of an individual (or genotype) as the measure of fitness. Theoretical results in the biological literature, however, revealed that the expected number of offspring generally is an inadequate predictor of evolutionary success in stochastic environments thereby asking for a modification of the propensity theory of fitness. In contrast, the relationship between expected trait frequencies and future trait frequencies has not received similar attention in the philosophical literature. In fact, the expected trait frequency is regularly invoked by philosophers of biology - either explicitly or tacitly - to predict future trait frequencies. The aim of this paper is to explore the relationship between the expected change in trait frequencies and future trait frequencies. I argue that inferring future frequencies from expected trait frequencies is generally flawed and discuss what inferences about future trait frequencies are warranted based on knowledge of the expected change in trait frequencies. More specifically, I distinguish between short-term and long-term predictions based on the expected change in trait frequencies from one generation to the next. While short-term predictions based on the expected change in trait frequencies are generally fallacious, I develop an argument warranting long-term predictions of trait frequencies based on the expected change in trait frequencies from one generation to the next. This positive result, however, is sensitive to the assumptions of the presupposed evolutionary model and as such reinforces the more general methodological lesson that one needs to be careful when inferring how evolution works based on specific mathematical models of the evolutionary process.

Alexander Aylward. Between Distance and Domination: Feminism, Experimentalism, and the Manipulation of Nature Who could deny that in order to produce knowledge about the world, we must engage with it in some way? Yet, philosophy of science was long dominated by theory-centrism, with little attention given to the realities of the experimental life—Alan Chalmers has called this the "theory-dependence tradition" (2013, 180). Experiments were merely tests of theoretical hypotheses, or else, "[o]ne simply began to philosophize on the assumption that science was capable of delivering a data base of settled observational statements" (Ackermann 1989, 185). Proponents of what Robert Ackermann terms "the New Experimentalism" (e.g. Ian Hacking, Allan Franklin, Peter Galison, David Gooding) have urged us to shift our philosophical attention to experimentation itself, and examine the material and embodied practices through which scientific knowledge is produced, in the laboratory and elsewhere. Following this lead, many contemporary perspectives on method, explanation, confirmation etc., are suffused with experimental talk—that of intervening in nature, of making changes, of manipulating.

For Ian Hacking, interfering with the world provides the most convincing evidence for the reality of certain unobservable entities (1983, chapters 11,16). Jim Woodward argues that "explanatory relations are the sorts of relations that in principle will support manipulations or interventions" (1997, 26), and has forwarded an account of causation based upon such manipulations and interventions (2003). Nancy Cartwright claims that the "kind of knowledge we have in modern science," is that which "provides us the understanding and the power to change the regularities around us and produce the laws we want" (1999, 78). Thus, the understanding of nature generated through scientific activity is that which services our meddling. Knowing the world is knowing how to alter it. Science, and knowledge-making generally, has been increasingly emphasised as a set of practices in which things are done and changes are made.

Whilst the notion of manipulating, intervening in, and thus exerting our material power over nature, has clearly gained currency with a great many philosophers of science, what I propose here to explore are this movement's feminist credentials. This paper asks how this recent trope of empiricism, with its emphasis on manipulating nature, should (or else should not) respond to or incorporate certain feminist insights.

As the first clause of my title suggests, feminists seemingly find themselves between the horns of a dilemma. Theory-centric philosophy of science distances the knower from the subject, thanks to the pervasive, metaphorical 'view from nowhere' (Nagel 1986). Along with feminist and postmodernist scholarship (Keller and Grontkowski 1983; Haraway 1988), the turn to practice in philosophy of science (e.g., Hacking 1983; Franklin 1986; Gooding, Pinch, and Schaffer 1989) speaks strongly against the ideal of scientists possessing abstract, objective vision. However, for several exponents of the new experimentalism, immersion in the material realm of experiment brings with it the occasion to dominate and manipulate aspects of nature—a trope oft disparaged by feminists (Merchant 1980; Keller 1985).

It seems the choice is between an abstract, theory-centric account of science in which material engagement between knower and known is neglected, and an account which recognises such engagement, but largely conceptualises it as dominative manipulation of nature. We are, then, stuck 'between distance and domination'. This paper—focusing mainly upon the work of leading new experimentalist Ian Hacking, his philosophy framed in terms of manipulating and intervening in natural phenomena—asks how feminists are to negotiate this turbulent landscape. Two historically rather independent literatures—that detailing and developing the new experimentalism, and that urging the appreciation of feminist insights in the practice and organisation of science—are here brought together, and their mutual bearings explored.

I review various possible resolutions of the 'distance-domination' dilemma, arguing that feminists have good reason to deplore the detachment of unmarked vision. Might feminists then evade the dilemma entirely by embracing the closeness of knower and subject, but eschewing all manipulative and dominative practices, instead exercising exclusively empathetic/interactive relations with nature? Whilst some early feminist critics of science seem to advocate as much, I maintain that such a move would be rash. A 'feminist science' need not commit unwaveringly to any particular account of the power-relation between inquirer and subject. Situation-to-situation, different relations are appropriate.

Marius Backmann. What's in a gold standard? In Defence of Randomised Clinical Trials The standardised Randomised Clinical Trial (RCT) is one of the most important and widely used experimental methods in use in medical research. It is a staple in evidence-based medicine, where it is often deemed the 'gold standard'of medical research. Above their application in medical research, RCTs have also been used in other fields such as Economics or practical policy making. Randomised Clinical Trials have been as popular as they are since they provide a standardised and easily applicable template that is meant to test claims about the efficacy of drugs and other substances, of economic measures, and of policies, in a controlled environment that is meant to eliminate statistically disruptive factors.

Recently, RCTs have faced significant criticism. Apart from practical worries e.g. concerning implicit bias in selecting test subjects, two major types of criticism can be differentiated. In my talk, I will first give a brief exposition of how RCTs work, where they are most commonly used, and how they are perceived as a gold standard. Following this, I will turn to the two fundamental types of criticism against RCTs and will argue against them.

The first type of criticism is a statistical worry: it is argued that we cannot possibly be certain that our sample is not atypical with regard to possible confounding factors. Due to the complexity of human organisms or, in the case of Economics and practical policy making, societies and economies, randomising does not guarantee that the sample is not skewed. John Worrall e.g. (Worrall 2002) argues that it is impossible to randomise all factors, confounding or not. As is the case with statistics, even if an outcome is unlikely, it might still be possible. So however unlikely, our test groups might just as well be skewed. But according to Worrall, even a more modest claim that randomisation makes it at least unlikely that the groups are skewed is problematic. The reason is that firstly, we do not re-run RCTs often enough to establish whether our grouping was atypical or not, and secondly, there is an infinite number of possible confounders. And if there are infinitely many possible confounders, then the possibility that our groups are skewed with regard to one of them might be high.

I will argue that at least in the case of medical research, we know enough about the relevant causal mechanisms in the body and about the mode of action of the tested substances from other sources than RCTs that we are justified to ignore a number of factors we have good reason not to expect to be disruptive. Worrall's line or argument ignores that we are not completely blind when it comes to eliminating confounding factors. If the Human body was a causal black box, then the argument would have force. However, we do e.g. know about the mode of action of new drugs before we test them in humans. Hence we are justified in believing that e.g. hair colour will not be a confounding factor, so we do not need to account for it. So the claim that the number of possible confounding factors is infinite is questionable. Moreover, Worrall's claim that we do not re-run RCT's often enough to rule out the possibility that our selection of test subjects might be unluckily skewed is also not necessarily true. A lot of studies are re-run over and over again, historically e.g. studies that dealt with the efficacy of homeopathic treatments.

The second type of criticism concerns what sort of claim RCTs really warrant. Nancy Cartwright e.g. (Cartwright 2007, 2010) argues that RCTs are used as what she calls „clinchers". She holds that, since RCTs are held to be the „gold standard", they should rigorously establish that within a certain subset of the population, a certain treatment causes a certain outcome. From this we deductively infer with the aid of a set of auxiliary hypotheses that the same treatment will also cause this outcome in the population. Cartwright holds that due to our imperfect sampling methods and the insecure nature of our knowledge of the similarity between the sample and the population, this last inference fails.

I will argue that RCTs should not be taken to deductively infer causal claims in the way Cartwright supposes. This seems to hold RCTs to a standard against which they must necessarily fail. It is no surprise that this deductive argument breaks down. But it seems a misrepresentation of scientific practice to reconstruct RCTs as a deductive method. I will argue that RCTs are an archetypical ampliative method and should not be reconstructed deductively. We do not know whether the distribution of confounding factors in sample and population is the same. And when we consciously select a sample that is more homogenised that the population, of course our inference that a treatment will cause a certain outcome in the population will be ampliative. But that it is ampliative cannot be an argument against RCTs: we know them to be fallible. But that doesn't entail that they're not the best we can do.

This is where the talk of a 'gold standard' becomes harmful: even in evidence-based medicine, , this method does almost never stand alone. But that medical research can make use of other sources of evidence like in vitro experiments to determine the mode of action for a certain substance does not entail that we could do away with RCTs. It would be uncharitable to read the admittedly imprecisely phrased handbooks and manuals for evidence-based medicine that praise RCTs as a gold standard to hold that they are a deductive method, that RCTs never go wrong in practice, or that no other method is indeed viable.

Tudor Baetu. On the Possibility of Designing Crucial Experiments in Biology The modern notion of crucial experiment emerged from the analysis of historical episodes where a single experiment seems to have conclusively and definitively sealed the fate of two or more competing hypotheses. The strategy behind such experiments hinges on the testing of the hypotheses under scrutiny relative to an aspect of empirical reality about which each of the competing parties makes a different prediction, such that the results will shift the balance in favor of the hypothesis making the correct prediction and against rivals that fail to do so. As philosophers of science like to point out, things are not quite as simple. The most famous challenge to a straightforward interpretation of the results of a crucial experiment is the underdetermination of scientific theory by evidence. One argument from underdetermination states that inferring that the theory making the wrong prediction must be false faces the problem of confirmation holism. A second argument from underdetermination states that even if it is possible to falsify the tested hypothesis, inferring that the theory whose prediction is confirmed must be true is questionable due to the problem of unconceived alternatives. A hypothesis cannot be confirmed against its rivals by means of crucial experiments because it cannot be ascertained that all possible alternatives have been considered. A similar shortcoming plagues abductive attempts to infer that the explanation that best responds to a set of epistemic virtues is true or the most likely to be true: the best explanation may simply be the best of a bad lot of false explanations.

In response to the challenge, it has been argued that the underdetermination thesis assumes a narrowly deductive view of explanation and confirmation, which is not descriptive of all science. In the context of different kinds of explanatory approaches, such as mechanistic explanations, the problem of underdetermination is thought to become more tractable because of the presence of constraints limiting the number of ways in which a certain phenomenon could be produced. Furthermore, Duhem described the reasoning strategy behind crucial experiments in terms of deductive elimination. There is no historical evidence to support this view. A more realistic rendering should emphasize the positive selection of the hypothesis supported by the experimental results instead of focusing exclusively on a negative selection strategy according to which the surviving hypothesis must be the correct one. For instance, if the hypothesis favored by the experimental results is also sufficient to explain the phenomenon without introducing additional assumptions, then we have no reasons to prefer a rival by itself incapable of accounting for the same results without adding further 'epicycles' to the explanatory story. Thus, a combination of evidence and simplicity considerations might favor a hypothesis against its rivals. With these arguments in hand, some authors–most notably Allan Franklin (2007), Sherrilyn Roush (2005) and Marcel Weber (2009)–proceed to argue that there are at least some unquestionable examples of successful crucial experiments in the history of science, of which the Meselson-Stahl experiment stands out as one of the most striking illustrations.

In this paper, I tackle the notion of crucial experiment from a different angle, by questioning the possibility of designing crucial experiments in the first place, leaving aside further epistemic difficulties brought about underdetermination issues. I argue that in some fields of investigation, most notably biological sciences relying primarily on mechanistic explanations, there are no sufficient reasons to believe that alternate explanations are mutually exclusive. This leads to an increase in the number of possible explanations, thus limiting the potential for designing and conducting crucial experiments.

I begin by analyzing the Meselson-Stahl experiment, showing that other experiments were required to support the semiconservative mechanism and argue that, given the significant efforts deployed in subsequent experiments, it is unlikely that the semiconservative hypothesis was accepted in virtue of being the best explanation. Moreover, contrary to arguments presented in the philosophical literature, the interpretation of the experimental results supporting the semiconservative hypothesis was not in any way simpler or less problematic than interpretations favoring rival hypotheses. Instead, I propose that the value of the Meselson-Stahl experiment lies in the fact that, in conjunction with additional experiments, it provided conclusive experimental evidence that E. coli DNA is replicated primarily via a semiconservative mode of replication. In turn, this finding turned out to be an important piece of the puzzle guiding the subsequent elucidation of the mechanism of DNA replication in prokaryotes, and then eukaryotes. I argue that instead of adopting a disjunctive elimination confirmation strategy, whereby relatively complete explanatory accounts are elaborated 'top-down' from substantive background knowledge and then pitted against each other by testing predictions about a particular aspect of empirical reality, science can also advance in a constructive conjunction manner, whereby individual pieces of experimental data about correlated and causally relevant factors are put together in an attempt to elucidate the mechanisms responsible for producing a phenomenon. The special case of the Meselson-Stahl experiment paves the way to a more general argument. While some fields of investigation benefit from background constraints that justify treating alternate explanations as mutually exclusive rivals, thus facilitating the pruning down of the total number of possible experimental outcomes to a handful of distinct scenarios each favoring only one hypothesis, other fields of investigation lack such constraints. If there are no reasons to believe that the proposed hypotheses are mutually exclusive, the number of possible experimental outcomes proliferates beyond control, leading to a situation where more than one hypothesis may predict the same outcome. Furthermore, if alternate explanations are not mutually exclusive rivals, it is no longer clear why evidence for a hypothesis should discredit other hypotheses and why lack of evidence for some hypotheses should count as evidence supporting the remaining hypotheses. In the conclusion of the paper, I summarize my findings and their implications for general philosophy of science.

Jonathan Bain. Emergence and Mechanism in the Fractional Quantum Hall Effect For some authors, an adequate notion of emergence must include an account of a mechanism by means of which emergent behavior is realized. These authors maintain that without such an account, emergence risks becoming a trivial concept that is appealed to whenever we lack epistemic access to a physical phenomenon, or the technical skill required to provide a complete description of it. According to Mainwood (2006, pg. 284), for instance, "...emergent properties are not a panacea, to be appealed to whenever we are puzzled by the properties of large systems. In each case, we must produce a detailed physical mechanism for emergence, which rigorously explains the qualitative difference that we see with the microphysical". The mechanism of most interest to Mainwood in the context of condensed matter physics is spontaneous symmetry breaking (SSB). Morrison (2012, pg. 160) similarly claims that emergence in condensed matter systems must be underwritten by a physical mechanism, and in particular SSB: "The important issue here is not just the elimination of irrelevant degrees of freedom; rather it is the existence or emergence of cooperative behavior and the nature of the order parameter (associated with symmetry breaking) that characterizes the different kinds of systems." Finally, Lancaster and Pexton (2015) note that while the fractional quantum Hall effect (FQHE) cannot be explained in terms of SSB, nevertheless a physical mechanism can be associated with it; namely, "long-range entanglement", and it is in terms of this mechanism that emergence in the FQHE should be understood.

The aim of this essay is to question this mechanism-centric view of emergence by considering Lancaster and Pexton's example of the FQHE in a bit more detail. The consensus among physicists is that this effect exhibits emergence, but there are at least four alternative explanations of it that, arguably, appeal to distinct ontological mechanisms, at both the microphysical level and the level of what have been called higher organizing principles. These explanations include (1) the Laughlin ground state account; (2) the composite fermion account; (3) the composite boson account, and (4) the topological order account. The FQHE is described by these accounts as (i) a many-body Coulomb effect of electrons, (ii) a one-body effect of composite fermions, (iii) a many-body effect of composite bosons, and (iv) a many-body entangled effect of electrons, respectively. These ontologically distinct microphysical mechanistic accounts are underwritten by the following ontologically distinct high-level mechanistic accounts: (a) localization (accounts 1 and 2); (b) spontaneous symmetry breaking (account 3), and (c) long-range entanglement (account 4).

In light of this underdetermination of mechanism, both microphysical and high-level, one is faced with the following options: (I) deny that emergence is present in the FQHE; (II) argue for the priority of one mechanistic explanation over the others; or (III) temper the desire for a mechanism-centric account of emergence. I will argue that there are good reasons to reject (I) and (II) and accept (III). In particular, I will suggest that emergence in the FQHE is best described in terms of what I will call a "law-centric" view of emergence. According to this view, emergence is characterized, in part, by novelty, and novelty is underwritten by an appeal to distinct laws, cashed out as the equations of motion associated with formally distinct Langrangian densities.

Daniel Bedingham and Owen Maroney. Thermodynamics and quantum information Since the earliest days of statistical mechanics, thermodynamic entropy has frequently been understood as a measure of our ignorance of the exact microstate of a macroscopically large system. The similarity of the mathematical expressions for the Gibbs entropy and the Shannon information measure have leant support to this view, to the point where Jaynes argued that entropy should be understood entirely in information theoretic terms [JAY1,JAY2]. However, simply having the same mathematical form should not be enough for this identification to be made: it must also be the case that a change in information should be accompanied by an equivalent change in entropy.

Landauer's Principle [LAN] seems to supply just the needed ingredient: it relates the change in the abstract information from a computation to a minimum thermodynamic cost in the form of heat generated in the environment. This minimum is defined purely in terms of the information processing operation itself and so every physical system which performs the computation must pay at least this cost. The cost is exactly given by the change in Shannon information over the computation, multiplied by a constant (Boltzmann's constant times ln 2) and the temperature of the environment, and is what would be expected if information and entropy were the same.

However, Landauer's Principle has only been properly studied in the context of classical information processing, built from logical operations such as AND, OR, and NOT gates. The development of quantum information and quantum computing changes our notions of information and its relationship with entropy. In quantum information processing, logical states are replaced by quantum signal states, and logical operations are replaced by quantum operations. A quantum operation is any allowed transformation that can be performed on a quantum state using the usual rules of quantum state development. This can involve auxiliary systems which are used to catalyse the operation by interacting with the signal states.

At first sight, the generalisation of Landauer's Principle to quantum information processing seems straightforward. The Gibbs entropy is replaced by the von Neumann entropy, and Shannon information is replaced by Schumacher information, maintaining the equivalence of the mathematical forms of information and entropy. It is relatively straightforward to derive an inequality which says the heat generated by a quantum operation is at least the change in Schumacher information over the operation, multiplied by a constant (Boltzmann's constant times ln 2) and the temperature of the environment [PAR].

However, an important part of the classical Landauer's Principle is that the relationship is a tight bound. While there are practical barriers to reaching the limit (such as finite size effects, time, etc.) there is no physical principle that prevents these barriers becoming arbitrarily small, and experimental tests are increasingly pushing at this boundary. When operating at the limit cost, the classical computation is being performed with thermodynamic reversibility.

We will use a simple argument to show that quantum computers cannot, in general, reach the limit for heat generation given by the Schumacher information measure, which is required for thermodynamic reversibility. This means that there is a necessary excess heat dissipation that is specific to quantum computation. Even in the limit of idealised heat baths, work reservoirs, and long times, the information theoretic cost given by the change in Schumacher information is not physically possible. Worse still for the attempt to forge a link between information and entropy, we will show that the minimum heat generation for a quantum operation cannot in general be expressed in terms of any function of the quantum information associated with that operation.

We will provide a rule to distinguish when a given operation can and when it cannot meet the condition of thermodynamic reversibility. We find this rule is satisfied for cases where thermodynamic reversibility is known to be possible: the classical limit of quantum computing in which the inputs correspond to orthogonal states; the case of pure unitary rotations; and resetting to a standard state.

Finally, in case all this may seem counterintuitive, we will discuss the paradigm of reversible computation, which is often suggested to solve the problem of heat generation in computing, and discuss how it applies to quantum computing. Quantum operations can always be implemented by unitary operations which act jointly on signal state and auxiliaries. These can always be run in reverse with no overall heat cost. The problem is that the reverse operation simply undoes the computation leaving no record of the output. In general the outputs of quantum operations will need to be recorded or passed on for further processing in a network of quantum computers. Reversible computing typically leaves junk auxiliary states that must be reset to be of further use. For classical information processing Bennett [BEN] was able to show that these auxiliary states could be reset at no cost, but to do so requires a copying operation that cannot be implemented for quantum states. In general the auxiliary states for quantum operations cannot be reset using Bennett's method, and they then carry the excess entropy cost, over and above the change in the Schumacher information associated with the operation.

At the quantum level, the tight connection between information and entropy appears to be lost.

Jonathan Birch. Shared Know-How Successful feats of cooperation manifest knowledge-how. When two people dance the quickstep together, without stepping on each other's toes or otherwise appearing badly coordinated, they manifest knowledge of how to dance the quickstep together. When two people row together in a coxless pair, moving smoothly through the water without veering from side to side, they manifest knowledge of how to row together. This observation leads naturally to the question: what is the relationship between the knowledge-how manifested by a group (e.g. a pair) of cooperating agents and the knowledge-how each agent individually possesses?

In broad outline, one can distinguish three possible approaches to making sense of this phenomenon, each corresponding to a well-known approach to shared intention. First, one might maintain that all the knowledge-how manifested in a joint action is, at bottom, individual knowledge-how. One might then seek to account for the shared knowledge-how of a group of agents in terms of the objects of, and relations between, their individual knowledge-how states. This approach would be modelled on Bratman's individual-centred approach to shared intention. Second, one might seek to account for shared knowledge-how in terms of a distinctive mode of knowing how: the idea, roughly, would be that an agent "we-knows how" to participate in a joint action, where "we-knows how" denotes some distinctive relation, subtly different from the familiar individual mode of knowing how. This approach would be modelled on those of Searle and Tuomela & Miller, who account for shared intention in terms of a distinctive mode of "we-intending", which may (for Tuomela & Miller) or may not (for Searle) be explicable in terms of individual intentions and beliefs. Third, and arguably most radically, one might locate the "sharedness" of shared knowledge-how in the existence of a plural subject: a subject that knows how to perform multi-agent actions, just as an individual knows how to perform single-agent actions. This approach would be modelled on Margaret Gilbert's "plural subject" theory.

Here I pursue the first approach: that is, I develop an account of the relationship between shared and individual knowledge-how in the spirit of Bratman's account of shared intention. For Bratman, shared intention does not require group subjects, or a distinctive "we-mode" of intending: instead, it arises from a subtle relational structure connecting the intentions of individual agents. The contention of this paper is that the same applies to shared knowledge-how: it too arises from a subtle relational structure, in this case connecting the individual knowledge-how states of the individual agents. My goal is to articulate the nature of the "subtle relational structure" in question.

Here is a brief outline. I begin by introducing three basic constraints that any adequate account must satisfy, which I call "Distribution", "Tether" and "Reliable success":

Distribution: If two agents S1 and S2 jointly know how to perform some shared cooperative activity J, it need not be true, of either S1 or S2, that he individually knows how to perform all the actions involved in a successful performance of J.

Tether: If neither S1 nor S2 knows how to perform any of the actions involved a successful performance of J, then S1 and S2 do not jointly know how to do J.

Reliable success: S1 and S2's jointly knowing how to do J explains why they are reliably successful at doing J (without aid from others) when they jointly intend to do J.

I proceed to consider a simple proposal on which, roughly speaking, two agents know jointly how to do J if and only if each individually knows how to do his part in J. I will call this the "simple distributive account":

Simple distributive account:

S1 and S2 jointly know how to do J if and only if, on forming the shared intention to do J, they respectively undertake individual intentional actions (j1 and j2) as parts of J such that:

i. S1 knows how to do j1;

ii. S2 knows how to do j2;

iii. The successful performance of j1 and j2 can, in principle, suffice for the successful performance of J.

This account is inadequate because it offers no explanation of the agents' reliable achievement of mutual coordination. This leads us to the key question: what is required for the reliable achievement of the form of mutual coordination that characterizes shared cooperative activities? I suggest that coordination in the context of shared cooperative activity consists in agents performing their parts in mutually coordination-enabling ways, while monitoring each other's performance and making responsive, success-promoting adjustments. These considerations lead us to an improved proposal, which I will call the "Mutual coordination account":

Mutual coordination account:

S1 and S2 jointly know how to do J if and only if, on forming the shared intention to do J, they respectively undertake individual intentional actions (j1 and j2) as parts of J such that:

i. S1 knows how to perform j1 in a coordination-enabling way while monitoring S2's (coordination-enabling) performance of j2 and making responsive, success-promoting adjustments his performance of j1.

ii. S2 knows how to perform j2 in a coordination-enabling way while monitoring S1's (coordination-enabling) performance of j1 and making responsive, success-promoting adjustments to his performance of j2.

iii. The successful performance of j1 and j2 can, in principle, suffice for the successful performance of J.

This account is the main positive proposal of the paper. I close by considering possible objections and important open questions. In particular, the account leaves open the question of how the mutual coordination account may help shed light on human social evolution. Michael Tomasello (2014) has argued for the importance of shared intention (or "joint intentionality") in underpinning uniquely human forms of cooperation. The suggestion, roughly speaking, is that shared intention was the critical evolutionary innovation that led to the dramatic divergence of the hominin evolutionary trajectory from that of the other great apes. My working hypothesis is that shared knowledge-how co-evolved with shared intention, and that the former provides a crucial but previously neglected component of the package of psychological adaptations that make human cooperation unique.

Daniel Burnston. Getting over Atomism Functional decomposition is one of our most important ways of understanding a biological system. Recently, the project has been wedded to that of mechanistic explanation—the attempt to explain biological phenomena as the result of types of causal interactions between distinct parts. When one understands how the parts of a system are causally organized, one understands how the phenomenon comes about.
Debates about the scope and limitations of mechanistic explanation, to a significant extent, have focused on whether functional decomposition is feasible. Functional decomposition is standardly construed atomistically. Atomism is the conjunction of two claims:
1. Each part is characterized by what it does intrinsically.
2. The function of the overall system is to be explained in virtue of the combined intrinsic functions of its parts, and not vice versa.
I argue, first, that both atomistic claims are false, and second that that this falsity has no bearing on the possibility of decomposition. My strategy will be to focus on three key properties—context-sensitivity, dynamics-dependence, and network-dependence—that have been taken to be incompatible with decomposition, and show that, once atomism is abandoned, these properties in fact support a robust notion of decomposition. I reconstrue decomposition as the ability to explain any given phenomenon in terms of specific interactions between parts, rather than as finding the intrinsic function performed by each part. I make the argument based on recent results in systems neuroscience. Areas of the brain are defined by their informational functions¬, namely what properties of a stimulus or task-situation they represent. Functional decomposition in neuroscience is the idea that one explains psychological behaviors by referring to the combined informational functions of the parts of the brain involved in producing them.
Context-sensitivity is a part P's performing distinct functions, conditional on what is going on external to it. These conditions can be either other conditions within the system or conditions in its environment, or (most often) both. Context-sensitivity is incompatible with intrinsic function, since what function P performs depends on factors outside of it. Arguments inferring failure of decomposition from context-sensitivity often rely on claims about the norms of explanation—if functions are context-sensitive, then we lack sufficiently generalizable explanations. I argue that this is false based on a comparison between perceptual areas in the brain. Area MT and V4, for instance, have significant overlaps in the types of information they process, particularly with regards to motion and depth. However, the contexts in which they process this information are distinct, and these distinctions in context allow the two areas to play non-overlapping roles in perceptual phenomena. Once we embrace context, and generalize over types of context, then there is no problem either with generalization or decomposition. This only defends the compatibility of decomposition with the falsity of claim (1), however.
Arguing that the falsity of (2) is compatible with decomposition involves answering the other worries in tandem. It is standardly thought that if dynamic interaction between parts, rather than serial causal interaction, underlies system behavior, then we cannot explain system behavior in terms of specific causal interactions between parts. Moreover, if the function of a given part P depends on what is going on elsewhere in the network, then we cannot explain the network in terms of what P does. Both of these claims, however, are false, as current work on neural coding shows.
Dynamics, far from being inherently contrary to decomposition, in fact helps implement specific functional interactions amongst otherwise contextually varying parts. A variety of recent studies have shown that neural signals are multiplexed. While the activity of any given cell or group of cells can carry multiple distinct types of information, certain properties of the signal—in particular, its frequency, can be used to disambiguate the multiple potential meanings. Put simply, even in a single spatio-temporal signal, different information can be encoded at distinct frequency bands. As such, a decoding system with sensitivity for the correct frequency can extract a particular sort of information from a signal whose overall semantic properties are multiple and contextually varying. An important mechanism for this selective readout is synchronized oscillation—one area receives a signal encoded at another by sharing a phase relationship with the sender. But if this is the case, then dynamic interaction precisely underlies functionally specific information transfer from one part to another, rather than denying its importance.
There is one more step to go, however. In a scheme such as this, something needs to determine what frequency an area will transmit at, as well as what frequency the receiver will decode. The way that this is implemented is through network-mediated synchronizing of brain areas. Studies in both working-memory and motor tasks have shown that different tasks implement different patterns of synchrony between distinct brain areas, and that individual cells in these brain areas represent distinct information depending on the frequencies of those oscillations. So—the representational function of individual cells and groups of cells within an area changes depending on network context. However, this does nothing to change the fact that, in a particular context, a single signal is output from a group of cells, which can then be decoded elsewhere in the network. Summed up: while network interactions help determine which context-sensitive function is implemented, in any given context a given area contributes a distinctive signal that helps explain task performance.
This last part is what is most important for decomposition. Say we wish to explain the difference in how the brain implements two tasks. Yes, we must refer to dynamics and network interactions, but we must also refer to the specific signals conveyed by specific parts of the brain. Different parts of the brain implement different patterns of contextual variation across distinct tasks, and in any given context, each part of the brain will contribute something specific to the behavioral phenomenon of interest. This is all that is needed for decomposition: if we abandon atomism's claims (1) and (2), we can decompose the brain functionally while taking into account both network effects and dynamics.

Fabio Ceravolo. Physically Unrestricted Composition Of the many challenges raised by Ladyman and Ross (2007) to the tenability of metaphysical propositions in the face of theoretical physics, the claim that physically informed answers to Van Inwagen`s (1991) special composition question will be highly disjunctive and sui generis has attracted very little attention.
The special composition question (SCQ) asks to single out circumstances that uniquely entail that two or more material objects compose. In the model that I wish to put forward, the answer to SCQ is neither disjunctive nor sui generis: It is necessary and sufficient for the composition of any two objects to occur at a world that the existence of a composite of such objects does not falsify the natural laws, the predictive consequences and the observations associated with the theory true at that world. I label this criterion law-sensitivity.
My strategy consists in showing that law-sensitivity arises from a different – and ultimately preferable – approach to metaphysical naturalism. Ladyman and Ross are correct in pointing out that, in physics, the conditions associated to the application of compositional predicates ('xs compose', 'y is a whole') are highly varied and mutually irreducible. However, that each application of a compositional predicate counts as a distinct answer to SCQ only follows if we have no other means of saying which things are composites aside from listing one by one the objects to which physics attributes compositional predicates.
I argue not only that we do in fact possess such means, but also that Ladyman`s and Ross` literal approach suffers from an independent problem.
The literal approach is inconsistent with unrestricted composition, the thesis that any two objects whatsoever compose, as indeed physics has no names for many sums of arbitrary parts. And this inconsistency seems ad hoc, for the attitudes of physicists and physics textbooks towards arbitrary sums tend to be agnostic rather than eliminative. For instance, when we consider whether there could be a sum of two far-distanced objects obeying Newtonian gravitation, say Mars and Alpha Centauri, what we expect to find is not a clear-cut negative judgement, but rather that the existence of the sum, for all that can be established by physical means, is compatible with the Newtonian laws, their consequences and the associated observations.
Sober (2015: ch. 1) calls the corresponding theoretical virtue, a non-eliminative variant of Ockham`s razor, the "razor of silence". Translating into the present context, the razor of silence applies to all cases in which, for a set P of natural laws, predictable consequences and direct observations, and for a set A of sentences stating the existence of a sum for every two objects, P entails neither the truth of all sentences in A ("everything whatsoever composes") nor the negation of some sentence in A ("some two things do not compose").
In consequence, I take it that literalism is best replaced by an approach that allows the SCQ to be settled by similar considerations of theoretical virtue.
In this new framework, the razor of silence opposes the eliminative version of Ockham's razor, which commands to eliminate (rather than remaining agnostic upon) 'overabundant' sums. I argue that the eliminative razor is best avoided, as its commands are unclear. Indeed, if the razor obliges one to eliminate all sums to which physics does not literally commit, then it will be inconsistent with the observation that physics recommends agnosticism rather than elimination. And if the razor commands the elimination of every object that is not strictly necessary for the (non-trivial) truth of physical laws, predictions and observations, then it is likely to eliminate all sums and lead to mereological nihilism.
Therefore, the best chance to combine agnostic judgements over arbitrary sums with a non-nihilist position is offered by a view that embraces all sums insofar as their existence is consistent with the background laws, predictions and observations. I take up this thesis and observe that it is neither sui generis, nor disjunctive.
Particularly, the account "updates" Lewis' (1986, 1991) view that it is suffices for sums to exist that they are ontologically innocent, whereby ontological innocence is now understood as consistency with the laws. Thus, all sums governable by the laws true at the world where composition occurs exist and are governed by such laws. All sums not governable by the laws true at the world where composition occurs either have no nomic behaviour whatsoever, or are governed by laws different from the laws of physics. Both possibilities suffice for triggering the eliminative razor.
In conclusion, I make use of two examples to illustrate a welcome consequence of law-sensitivity. Namely, the difference between unrestricted and restricted composition comes out theory-relative, in that whether every two objects whatsoever or only some objects compose is a function of the particular laws true at the world where composition occurs.
In Newtonian mechanics, consistency with the laws demands that pairs of Newtonian masses have a centre of mass located at an averaged distance (determined by mass distribution), which feels the vector addition of the forces felt at the centres of individual masses. These conditions obtain for every pair of Newtonian masses. Hence sums of every two Newtonian masses are governable by the Newtonian laws: unrestricted composition holds in a Newtonian world.
In quantum chromodynamics, the laws regulating colour interactions state that all and only quarks at distances smaller than the hadron diameter instantiate the mutually attractive colour force and that only quarks instantiating the colour force are observed to promote quark-antiquark pairs from the vacuum when enough energy is supplemented. This couple of conditions does not obtain for every two quarks whatsoever. Pairs of far-distanced quarks, for example, do not promote quark-antiquark pairs from the vacuum when enough energy is supplemented: their sum cannot be governed by the quark-laws – the eliminative razor applies. Hence restricted composition holds in a quark-world.

Dimitri Coelho Mollo. Deflating Representational Content Representation plays a central explanatory role in the cognitive sciences. In order to fulfil this role, theories of representation must meet some requirements: explain how representations come to have the contents they do; give an account of what makes them into representations; and make space for the possibility of misrepresentation. Preferentially, all these requirements should be met by having recourse exclusively to naturalistically acceptable entities and relations, so as to give the notion of representation a respectable place in the scientific worldview.
The notion of representation has traditionally come hand-in-hand with that of computation. The idea that the cognitive system is to be viewed as computational is one of the founding pillars of cognitive science. Explaining what computing systems are and in what way the cognitive system is computational, or usefully explained as so being, is another foundational issue in the cognitive sciences.
In this paper, my aim is to provide the outline of a theory of representation able to play the required explanatory role while steering clear from the metaphysical difficulties that plague existing theories. My approach will be deflationary. I will rely on the robust notion of concrete computation provided by the mechanistic view (Piccinini, 2015; Milkowski, 2013) to individuate computational structure as one of the factors that carries the most load in explaining complex appropriate behaviour. Ascription of determinate representational content comes on top of that, and heavily depends on the task at hand and on the particular situation the organism finds itself in.
I will be accepting the invitation made by Piccinini (2004) to conjoin existing theories of content, which have traditionally relied on a semantic view of compu-

tation (Fodor, 1975; Shagrir, 2006; Sprevak, 2010), with the non-semantic view of computation provided by the mechanistic account. I take that Structural Representation (Swoyer, 1991; Cummins, 1996; Ramsey, 2007) is a particularly promising candidate for such a treatment. It is a notion of representation which successfully answers Ramsey's 'job description challenge' and which is often at work in empirical research.

Structural Representation is based on the idea that representations represent what they do by virtue of instantiating the same relational structure, i.e. by being in a second-order resemblance relation, to what they represent. Maps are the clearest examples of (non-mental) representations that structurally resemble what they represent. In a city map, physical (spatial) relations (distance of points in the map) mirror the spatial relations of things in the world (streets, buildings, squares, etc.).

Structural Representation has as its nemesis the problem of liberality --- given that structural resemblance is a liberal relation, any representation will represent many different things, leading to wild non-uniqueness of content (Goodman, 1976; Shea, 2013). This, on its turn, hurts the explanatory purchase of the notion of content.

I suggest that the robust notion of computation and thereby of computational structure that the mechanistic view of computation provides can come to complement the notion of Structural Representation and make it a better candidate for a theory of representation.

Viewing the cognitive system as a computational mechanism allows the non-semantic individuation of its relevant functional structure. The elements and processes of the system are carved up according to their computational roles. As such, the cognitive system has an objective computational structure, an organisation of its elements and processes that play a computational role and stand in computational relations to each other.

Computational structure mechanistically-individuated provides a promising way of cashing out the relevant relational structure of representational vehicles in Structural Representation. A representation would thus represent all the entities in the world that share its computational structure. This introduces strong constraints on what structures of the cognitive system are candidates for representational status, which helps to curb the liberality of the account. Nevertheless, some liberality is still present. Computational structures will bear resemblance relations to many different target domains: representations will still have non-unique content.

I recommend that we get around this problem by metaphysically deflating the notion of representation. Given the robust non-semantic account of computation that the mechanistic view offers, it becomes possible to deflate the notion of representation while keeping to realism about representational vehicles, i.e. computational structures.

The computational structure of internal states and processes lies at the basis of ascriptions of representational content. Content is to be seen as what captures the successful use of an internal state in the context of certain task domains, in which mechanistically individuated computational structure plays a central role.

Partially shifting the explanatory burden to computational structure allows us to lift much of the weight traditional theories of cognition place on the notion of representational content. Representational content is not needed to individuate the relevant explanatory states, nor is it required to avoid the trivialisation of computational implementation. This clears the way for a deflated notion of representation. On the metaphysical footing of representational content, two deflationary paths worth investigating open up.

One is pragmatism about representational content (Egan, 2013). According to this view, the notion of representational content grabs a hold only in light of our explanatory interests. Representation is invoked relative to specific scientific projects and to measures of behavioural success dictated by the interests of theorists.

Another is a sort of 'mild realism', inspired by Dennett (1991). The basic idea is that representational content captures certain patterns in nature: the complicated regularities involving the interactions between organism and world --- across contexts, despite disturbing conditions, and so on. Content integrates disparate and apparently independent contributions in a whole that is explanatorily fruitful.

I remain largely neutral on which of these two metaphysical views to adopt. For all scientific purposes the distinction between them is moot, and there may be no compelling reason to choose one over the other. The account, in its two possible declinations, is not merely instrumentalist, given the reliance on computational structures and computational mechanisms. For the same reason, neither is it eliminativist. The place of representation in the cognitive sciences is preserved, unencumbered by metaphysical burdens.

Erik Curiel. On the Cogency of Quantum Field Theory on Curved Spacetime Quantum field theory on curved spacetime (QFT-CST), postulates that it makes sense in certain regimes to treat the geometry of spacetime as classical while treating the matter that geometry couples with as quantum fields. The form of that coupling is defined by the semi-classical Einstein field equation (SCEFE), equating a classical geometrical structure, the Einstein tensor, with the expectation value of the stress-energy tensor (considered as a quantum operator). There are, however, many serious, unresolved technical and conceptual problems with this framework about even such basic issues as its physical consistency. In this talk, I plan to canvass those problems and briefly discuss how serious they are. Some of the problems I discuss are known to physicists, but, though they are deep conceptual problems in one of the most active and fundamental branches of theoretical physics today, including black-hole thermodynamics and early-state cosmology, they have not been addressed in the philosophy literature. Other problems I plan to discuss do not seem to be considered even by physicists. The problems fall into two classes: those with the consistency of the SCEFE itself; and those arising from the difficulty of formulating analogues to the standard energy conditions of general relativity (GR).

In the first class, one possible inconsistency in QFT-CST I have not seen discussed is that, in GR, it is not always clear what, if any, the physically significant differences are between "matter" and "gravitation". Mathematically, the difference between the two is partially captured by the difference between Ricci and Weyl curvature. (There is non-trivial Ricci curvature only where there is "matter".) One may therefore ask with some justice, why consider quantum effects in the Riemann tensor associated only with the former type of curvature but not the latter? Can one even consistently define a Riemann tensor that is "part quantum, part classical"?

There is an even deeper problem here, however. In general, there are severe technical difficulties in trying to define the quantum operator representing classical stress-energy tensor, even in Minkowski spacetime. The problems in curved spacetime are even more difficult, arising from the generic lack of a timelike Killing field. Even if one admits the possibility of the construction of such an object, however, there are still difficulties facing the attempt to define a reasonable notion of its expectation value. Those problems go even deeper than the lack of a timelike Killing field, rather arising from the lack of an unambiguous notion of parallel transport of vectors (due to curvature), leading to ambiguity in the standard calculational techniques for computing an operator's expectation value.

Even if one were to solve all these technical and conceptual difficulties, or at least argue convincingly that they can be set aside at least for the sorts of practical purposes that physicists have, there remain problems that bedevil even those practical purposes. Because the SCEFE has the same character as the back-reaction equation for classical charged particles, one must deal with the problem of divergent solutions and their associated pathologies. Although some physicists have argued with some plausibility that, to some degree, this problem can be mitigated by restricting attention to perturbations off flat spacetime, there remains the fact that it is parlous at best to try to extract physically significant, generic predictions from the SCEFE.

Now, for the second class of problems. The standard energy conditions play a central, fundamental role in GR: as assumptions in essentially all of the deepest and farthest-reaching results (e.g., all singularity theorems and the Laws of Black-Hole Mechanics); and their failure allows for every kind of pathological behavior (e.g, closed timelike curves and naked singularities). The status and physical interpretation of those energy conditions is still an open problem even in classical GR. They are yet more problematic in QFT-CST spacetime. First, it is not clear even how to formulate them, given the problems in representing stress-energy in a way that respects both the quantum nature of the fields and the classical nature of the underlying geometry. Second, once one has fixed any of the known formulations, it is almost ridiculously easy to construct physically reasonable generic violations of it.

There have been several attempts to address these problems. The most widely used of the resulting formulations, and the one that physicists seemingly have the most faith in, is the so-called Semi-Classical Average Null Energy Condition (SCANEC). Early efforts to prove that the SCANEC in classical GR is satisfied by "physically reasonable" fields seemed promising. It was quickly realized, however, that it is generically violated by reasonable matter fields, even when the spacetime is "nearly flat". Even worse for our purposes, it has been shows that even "physically reasonable" solutions to the SCEFE, when they can be constructed, generically violate the SCANEC. Some physicists argue that such violations are generically restricted to spacetime regions at the Planck scale, where one expects the semi-classical model to already have broken down, concluding that any pathologies associated with violations of the SCANEC (e.g., traversable wormholes and closed timelike curves) should be restricted to the Planck regime. There are serious technical and conceptual problems with such arguments, however, that make those conclusions suspect: many of the calculational techniques are physically unmotivated; and the calculations use only linear (i.e., non-self interacting) quantum fields. The case of non-linear quantum fields is so different from that as to make it more of a pious hope than anything else that one can draw any relevant conclusions to that case from the linear one.

All the problems I discuss ramify into essentially every philosophically important question surrounding the relationship between quantum physics and gravitational phenomena, including: whether QFT and GR are necessarily inconsistent; whether the semi-classical approximation of QFT-CST is physically well motivated and, if so, what the proper interpretation of its results are; and whether or how the effects of QFT-CST (e.g., Hawking radiation) can give insight into a possible theory of quantum gravity.

Chloé de Canson. Probability in Everettian Quantum Mechanics This paper is concerned with the so-called consistency problem of probability in Everettian Quantum Mechanics (EQM): what sense can we make of the probabilities postulated by the Born rule given that EQM is deterministic and that every outcome of a given measurement actually obtains? Section I underlines two related major issues that arise in most of the proposals aiming at making sense of the uncertainty associated with probability. Section II relies on the distinction between the splitting and diverging pictures of EQM to explain why these proposals face such issues, and to propose a working alternative. Finally, Section III argues that the understanding of EQM advocated in Section II does away with objective chance altogether. The probabilities postulated by the Born rule are neither chance nor credences, they are a third type of probability which the paper calls descriptive probabilities, and which are argued to be a suitable interpretation of probability for a much wider range of phenomena beyond EQM.

I. Probability and Uncertainty

For a phenomenon to be explained in terms of genuine probabilities, it is argued, there needs to be uncertainty about the outcome. But how can there be any uncertainty associated with an experiment in EQM? The observer knows that all outcomes will obtain! Those arguing that there is uncertainty concede this, but nonetheless claim that the agent is uncertain about which of the outcomes she will observe. The intuitive force behind this argument lies in the fact that we only observe one outcome in a given experiment. There are however two main problems with all versions of such arguments: the first one to do with personal identity, and the second to do with propositional attitude reports. The paper outlines these problems, highlighting how the difficulties arise in EQM. The paper then considers an alternative way of understanding uncertainty, originating from Vaidman, but constructs a thought experiment to show that, contrary to what Vaidman claims, his account collapses to the previously considered ones and therefore encounters the same problems. The sections ends on the conclusion that all such attempts to construct uncertainty in EQM are at best problematic, and if an alternative way of introducing uncertainty can be found which evades these problems, then that would strongly count in its favour.

II. Splitting and Divergence

Saunders (2010) proposes a distinction between what the paper calls the splitting picture and the diverging picture of branches in EQM. On the splitting picture, branches have a segment in common, which splits at measurement. On the diverging picture however, the branches are at no stage numerically identical, but are qualitatively identical prior to, and different post, measurement. Saunders shows that either picture provides an adequate Everettian understanding of the quantum formalism, which leads Wilson (2012) to claim that the choice between the two should be made on the basis of coherence and theoretical utility. The distinction enables the present paper to explain why the attempts at making sense of uncertainty considered in the previous section run into serious problems -- this is a consequence of having (unknowingly?) adopted the splitting picture. Because, as it is shown, it is easy to make sense of uncertainty

in the diverging picture, this is the one, the paper argues, that should be adopted.
III. Non-Chancy Objective Probabilities
The final section investigates the consequences of adopting the diverging picture for our interpretation of probability in EQM. More specifically, it argues that, because it is deterministic, the diverging picture is inconsistent with objective chance. It argues however that the Born rule, if it is understood as giving the mod-squared amplitudes of branches, is an objective feature of the universe; and therefore must be understood in terms of objective probabilities. A non-chancy objective interpretation of probability, called descriptive probabilities, is presented, and is shown to explain both how objective probabilities enter EQM and how uncertainty is preserved via a credence function.
This last section is centred around the debate of the compatibility of chance and determinism. The problem is usually phrased in terms of classical statistical mechanics (CSM), a deterministic theory that is widely held to postulate non-trivial objective probabilities. The paper however redirects these arguments in the case of diverging EQM, which is arguably similar to CSM in all relevant respects. It particularly engages with arguments by Loewer (2001), Schaffer (2007), and Lyon (2011). The overall argumentative strategy is the following. A conceptual analysis of what is usually meant by 'chance' yields that it is incompatible with determinism. The reason why that has been commonly rejected by compatibilists is because they rest on a false dichotomy between objective chance and subjective credence. But, there are in fact more than one interpretation of objective probability: a probability function might be well-suited to objectively describe a system. For example, it is an objective fact that 20% of my fingers are thumbs. When an agent says that 20% of her fingers are thumbs, she is making a meaningful claim which involves probabilities, but which does not involve chance! The paper argues that descriptive probability is the correct interpretation of probability to understand the Born rule. It then devises a principle similar to David Lewis' (1980) Principal Principle, which says that if an agent knows a descriptive probability function, then her credences in the events on which that function is defined should match the values postulated by the function. So, if an agent knows that 20% of her fingers are thumbs, her credence in that any given one of her fingers is a thumb should be 20%. This is shown to simply be a probabilistic version of Jeffrey's claim that agents should conditionalise on all the available evidence. It is also shown to imply the desired uncertainty in diverging EQM. The paper concludes by claiming that, given that chance is notoriously difficult to make sense of, the fact that diverging EQM lets us do away with it should be taken to be a strong argument in favour of the Everett interpretation.

Sebastian de Haro. Diffeomorphism Invariance versus Duality, or: The Holographic Hole Gauge/gravity duality relates a theory of gravity in a (d+1)-dimensional spacetime to a d-dimensional conformal field theory (CFT) on the conformal boundary of the spacetime. An important philosophical question is: how are diffeomorphism invariance and duality related? Does this relationship mesh well with extant philosophical doctrines about gauge invariance?
The general expectation is that gauge symmetry and duality trivially commute with each other, because duality is supposed to relate only quantities and states that are gauge invariant (in this case: diffeomorphism invariant). This may be so in general; but in this talk I will give a result about gauge/gravity dualities that indicates that two dual theories are 'closer in content' than you might think. For each of an important class of gauge symmetries in the gravity theory (diffeomorphism symmetry) is mapped by the duality to a symmetry of the CFT. This is worth stressing since some discussions suggest that all gauge symmetries in the bulk theory will not map across the boundary theory, but instead be 'invisible' to it. As we will see, this result also prompts a comparison with the Hole Argument. And there is an interesting interpretive fork for dualities, that bears on the extent to which diffeomorphisms are cases of gauge symmetry.
The above discussion thus prompts two urgent questions:
(i) Can one characterise the class of (d+1)-dimensional diffeomorphisms which gives rise to the spacetime symmetries of the CFT?
(ii) Are the remaining diffeomorphisms 'invisible' to the dual CFT?
In the physics literature, the first question has been answered in some specific cases (Imbimbo et al., 2000). On the second question, a heuristic notion of 'invisibility' has been introduced in order to argue that the diffeomorphisms of the gravity theory are not 'seen' by the QFT at all, and that they are 'emergent' (Horowitz and Polchinski, 2006).
In this talk I will discuss results from conformal geometry that bear on both questions above and correct some of the heuristic expectations. In order to disentangle different aspects of the notions of 'visibility' and 'invisibility', I will introduce four relevant conditions on a diffeomorphism F, from which I will construe a sharp notion of 'invisibility':
(1) F preserves the 'normal form' of the (d+1)-dimensional metric.
(2) F tends to the unit map near the boundary.
(3) F is an isometry of the d-dimensional conformal manifold.
(4) F leaves all QFT correlation functions invariant.
For my construal of visibility and invisibility of a diffeomorphism, I will rely on a theorem (which I prove in my paper, building on Fefferman and Graham (1985, 2010)), in simplified form:
Theorem. Let F:U-->U be a diffeomorphism in an open neighbourhood of the boundary of a (d+1)-dimensional manifold with conformal boundary. Under very general and natural assumptions, if the metric satisfies Einstein's equations in vacuum with a negative cosmological constant, then the following statements hold:
(Visible) If F satisfies conditions (i) and (iii) above or, alternately, if it satisfies conditions (i) and (ii), then it reduces to a conformal transformation on the boundary. We will call such diffeomorphisms visible, because the CFT fields transform non-trivially under them.
(Invisible) There exist non-trivial diffeomorphisms F satisfying conditions (ii), (iii), (iv). Such diffeomorphisms will be called invisible.
(Trivial) If F satisfies conditions (i), (ii), (iii), then F is the identity map. Thus, there are no invisible diffeomorphisms that also satisfy condition (i).
This theorem thus characterises a class of diffeomorphisms that are 'visible' and correspond precisely to the spacetime symmetries of the CFT. It also identifies a class of diffeomorphisms that are 'invisible' to the CFT, hence are genuine diffeomorphisms of the gravity theory which are not mapped by the duality. I will discuss the sense in which these diffeomorphisms are 'emergent', and how their emergence relates to extant accounts of emergence in the literature on dualities.
The theorem also has implications for Einstein's infamous Hole Argument: for the class of diffeomorphisms that can be used for the argument turns out to be smaller than expected. The theorem means that, for a certain class of non-isometric, diffeomorphically-related scenarios, demanding that the diffeomorphisms be 'invisible', as in (i)-(iii), is sufficient to show that they are also trivial. The details of the Hole Argument are thus constrained in a particularly interesting way, when studied in the context of gauge/gravity duality. But the theorem also shows that one does not need to go to full gauge/gravity duality in order to obtain these results—they already in some sense exist within the realm of classical conformal geometry.

Finnur Dellsen. What Scientific Disagreement Tells Us About Rational Consensus Most people who are not themselves scientists do not have access to scientific evidence or the expertise to analyze such evidence. Instead, they evaluate scientific theories indirectly by relying on the testimony of those who are experts on the relevant topics. However, one salient fact about scientific experts is that they frequently reach conflicting conclusions, often on the basis of the same evidence. In public discourse, this fact is often taken as a reason not to trust such experts, even on matters on which there is little or no disagreement.
In direct opposition to this common view, I argue here that the very fact that there is disagreement among experts on a given issue provides a positive reason for us to trust those experts concerning theories on which there is consensus. I show how this view can be argued for in three distinct (but in my view compatible) epistemological approaches: a Bayesian approach, an Explanationist approach, and a Robustness approach. Each argument exploits the same basic idea, viz. that the fact that a group of experts frequently disagree suggests that they wouldn't reach a consensus on a theory unless there really are strong epistemic reasons for believing it.
1. The Bayesian approach
Let Agr(H1) be the proposition that the experts on the relevant scientific topic agree that H1 is true; let Dis(H2,…,Hn) be the proposition that experts disagree about H2,…,Hn; and let J(H1) be the proposition that H1 really is epistemically justified by the scientific evidence available to the experts. Let us start by noting that since the likelihood of J(H1) on Agr(H1) is presumably greater than the likelihood of ~J(H1) on Agr(H1) – i.e. since Pr(Agr(H1)|J(H1)) > Pr(Agr(H1)|~J(H1)) – we have that Agr(H1) by itself (incrementally) confirms the probability of J(H1). In other words, expert consensus on H1 confirms the proposition that H1 is justified since experts are likelier to reach consensus on H1 if H1 is justified than if H1 is unjustified. This should not be surprising.
What might be surprising is that Dis(H2,…,Hn) boosts the confirmation of J(H1) provided by Agr(H1). To see why, we first note that it is easily provable that Agr(H1)&Dis(H2,…,Hn) confirms J(H1) to a greater extent than Agr(H1) by itself iff Pr(Dis(H2,…,Hn)|Agr(H1)&J(H1)) > Pr(Dis(H2,…,Hn)|Agr(H1)&~J(H1)). This inequality arguably holds since experts who have reached a consensus on an unjustified hypothesis would presumably be more likely (than experts who have reached a consensus on a justified hypothesis) to also reach a consensus on other hypotheses regardless of whether those hypotheses are justified. If that's right, it follows that expert disagreement on H2,…,Hn adds to the confirmation of J(H1) beyond what is already provided by the expert consensus on H1.
2. The Explanationist approach
According to Inference to the Best Explanation (IBE), inferring H from E is warranted if H provides a better explanation of E than any available rival explanatory hypothesis. Note that any plausible version of IBE presupposes a "requirement of total evidence" to the effect that E must include all of the agent's (relevant) evidence. This means that whenever experts both agree on H1 and disagree on H2,…,Hn, an explanatory hypothesis inferred by IBE must be judged according to its capacity to explain both the consensus on H1 and thedisagreement on H2,…,Hn.
Now consider two rival explanations for Agr(H1): (a) [CritEval]: Scientific experts form beliefs by examining the scientific evidence for a given hypothesis critically and/or independently of each other; accordingly, they reached a consensus on H1 because H1 was overwhelmingly supported by such evidence. (b) [CrowdPsych]: Scientific experts form beliefs irrespective of the evidence by following the lead of their peers and/or the lead of some scientific authority; accordingly, they reached a consensus on H1 because H1 was believed by the relevant peers and/or authority.
When Agr(H1) is one's only relevant evidence, these explanations would arguably be on a par. However, the situation changes if Dis(H2,…,Hn) is also part of one's evidence. To see why, note that [CritEval] can easily explain Dis(H2,…,Hn) as due to the scientific evidence for H2,…,Hn not being sufficiently univocal for scientists to reach identical conclusions. [CrowdPsych], by contrast, would have to invoke some special reasons why scientists would not also follow each other's leads (or that of an authority) on H2,…,Hn as they are alleged to have done regarding H1. Such an explanation would suffer either in simplicity (if it invokes such special reasons) or explanatory scope (if it does not). The upshot is thus that disagreement on H2,…,Hn helps make [CritEval] a better explanation than [CrowdPsych]. Since J(H1) follows from the former but not the latter, this helps make J(H1) warranted by IBE.
3. The Robustness approach
A result is robust just in case it holds under a variety of different assumptions about a given phenomenon. Since a result's robustness suggests that it is not an artifact of any of the specific assumptions that are being made about the phenomenon, robustness (fallibly) indicates that the robust result really holds true of the phenomenon. It is worth emphasizing that it is crucial to robustness that there is genuine variety in the relevant assumptions – a result that holds only under very similar assumptions would not be significantly robust. We now apply this to the consensus-disagreement situations in which we are interested.
The fact that experts disagree on H2,…,Hn indicates that they have a variety of ways of evaluating evidence in their area of expertise, e.g. that they have different background assumptions and/or different standards for what counts as good evidence. And the fact that these experts still agree on H1 indicates that H1 would be justified according to any of these ways of evaluating the evidence. Thus, the combination of consensus on H1 and disagreement on H2,…,Hn indicates that H1's justification is robust in the relevant sense, and hence that H1 really is justified. Here, as before, disagreement on H2,…,Hn is shown to indicate that a consensus hypothesis H1 is epistemically justified.

Joe Dewhurst. From Folk Psychology to Cognitive Ontology The conceptual categories provided to us by folk psychological discourse are typically taken for granted, with the notable exception of the arguments given by the various eliminative materialists (see e.g. Feyerabend 1963; Rorty 1965; Churchland 1979, 1981; Stich 1983). In philosophy this is manifested in the way we talk about mental states, the examples we choose when constructing thought experiments, and the heavy emphasis on the folk psychological concept of belief in epistemology. In cognitive science the impact is subtler but nonetheless important, manifested in the phenomena that researchers choose to study and the terminology that they use to report their results.

Our reliance on folk psychological concepts poses a risk to both philosophy and cognitive science. This is because, quite aside from any concerns about elimination as such, we have good reason to think that folk psychological concepts lack two features that would make them suitable for philosophical and scientific application. These are a.) stability across cultures and contexts, and b.) sufficient fineness of grain. Without these two features folk psychological concepts become liable to distort the validity of arguments given in philosophy and of explanations given in cognitive science.

The first feature, stability across cultures and contexts, is important if philosophical and scientific results are going to be generalisable and repeatable. It is also important if we want to be confident in taking our own folk psychological intuitions to be reflective of intuitions that are held pretty much universally. If folk psychology does not exhibit this feature then we should not expect concepts drawn from our folk psychological discourse to be either universal or generalisable, which will consequently impact the reliability of any theory or experiment that makes use of such concepts.

We have evidence that folk psychology is not stable across cultures and contexts. In a comprehensive survey of ethnopsychological research, Lillard (1998) draws attention to both differences of emphasis in folk psychological discourse, as well as explicitly divergent conceptual taxonomies. More recently, a collection of short position papers from an interdisciplinary conference titled "Toward An Anthropological Theory of Mind" (Luhrmann 2011) gives a flavour of the kinds of cultural variation that can be found in folk psychological discourse.

The second feature, sufficient fineness of grain, is important if philosophical and scientific characterisations of cognitive systems are going to be able to capture the full complexity of those systems. This is not to say that we should endorse a naïve reductionism where any mental state must be reduced to the finest grain of microphysical state available to us, but rather that any characterisation of a cognitive system must be sufficiently fine grained to capture any functional distinctions that are relevant to the explanatory project at hand. If folk psychological concepts are not able to do this, then we must replace them with more appropriate concepts, or else risk oversimplifying the systems that we are trying to study.

Many folk psychological concepts appear to struggle when applied in scientific contexts. Examples include the folk concept of belief in social cognition (De Bruin & Newen 2014: 303) and elsewhere (Gendler 2008), the folk concept of memory (Danziger 2008), folk taxonomies of the senses (Macpherson 2011), and the (European-American) folk concept of consciousness (Irvine 2012).

Whilst classical eliminative materialists have argued that folk psychological discourse should eventually be replaced by a more scientifically respectable taxonomy, this is not a necessary (or even desirable) consequence of acknowledging that folk psychological concepts might not be well suited to technical applications in philosophy and cognitive science. Even if folk psychological concepts turn out not to exhibit either of the two features introduced above, folk psychological discourse might nonetheless continue to serve a number of important social functions, such as facilitating a narrative understanding of human behaviour and enabling certain kinds of normative and pedagogical discourse (McGeer 2007; Zawidzki 2013).

This leaves us with the question of what to replace our folk psychological concepts with, if we are no longer permitted to use them in philosophy and cognitive science. This paper will propose that we should adopt a systematic methodology for gradually revising folk psychological concepts in light of cognitive scientific evidence, building on recent work in the field of cognitive ontology (e.g. Poldrack 2010; Anderson 2015). This will involve triangulating evidence from multiple disciplines including anthropology, psychology, and cognitive neuroscience, and using this evidence to implement iterative improvements to experimental designs and theoretical models (cf. Turner 2012). The resulting position will be partially eliminativist with regard to the scientific and philosophical application of folk psychological concepts, but respectful of the positive role that folk psychological discourse might play outside of academic contexts.

Anna Maria Dieli. Criteria for cell individuality: cancer from a multi-level selection perspective Cancer, also known as malignant tumour, is a disease involving aberrant proliferation of cells and the ability to invade other tissues. In cancer, cells grow out of control and become invasive: therefore, it is usually described as a cell disease. Mutation, competition and natural selection between cells are thus the main components of the phenomenon of cancer (Nowell 1976). Mutations that arise in cancer cells give them a selective advantage on normal cells: therefore, cancer cell proliferates as a successful Darwinian lineage. Cancer cells may thus be described as a Darwinian population subject to natural selection. In this framework, cancer cells fulfil the criteria for Darwinian evolution by natural selection, which is heritable variation in fitness: investigating cancer in a Darwinian perspective has generated new insights into disease aetiology, pathogenesis and treatment.

This talk aims to analyse whether the Darwinian framework is useful to understand cancer cell identity. It will be maintained that it is correct – to some extent – to describe cancer cells as Darwinian individuals; notwithstanding, cancer cells identity cannot be understood through this framework. In fact, cancer phenomenon cannot be described merely from a cellular point of view. A cancer cell has to progress into a normal tissue in order to be considered as pathological. When putted "in vitro", a cancer cell is not distinguishable from any other cell: it is just a cell which replicates. For instance, it has been proved that transplanting a cancer cell in a normal tissue not always gives rise to a tumour. A tumour arises from the interaction between cells, tissues, organs and the whole organism.

In the first part of the talk, the paper by Germain (2012) will be analysed, in order to inquire whether cancer cells satisfy the formal requirements for being Darwinian individuals (Germain 2012). According to Godfrey-Smith (2009), Darwinian individuals are identified thanks to a list of characteristics: fidelity of heredity, abundance of variation, continuity of the fitness landscape, dependence of reproductive differences on intrinsic character, reproductive specialization and integration. If a cancer cell population has all these characteristics – at least to some degrees – it can be considered as a Darwinian population. The aim of this author, therefore, is to understand "how Darwinian" cancer cells are. In cancer cells, individual differences in fitness are less dependent on outside signals and more dependent on intrinsic features. This is why a cancer cell cannot be considered to be a paradigmatic Darwinian individual: it is re-darwinized, that is, selection acts again at cellular level. Cancer cells are no more subjected to the organism constraints. Natural selection acts on cancer cells as it acts on autonomous entities in nature; and the action of natural selection at cellular level destroys the integration of the organism. Therefore, because of the strong dependence of fitness differences on intrinsic characters, cancer cells cannot be considered to be paradigmatic Darwinian individuals.

In the second part of the talk, the paper by Lean-Plutynski (2015) will be analysed, in order to understand why cancer can be described as a multi-level selection phenomenon. According to these authors, cancer is both a subject to selection at multiple levels and a by-product. First of all, cancer is an example of multi-level selection. In multilevel selection, selection acts at more then one level simultaneously. Multilevel selection can be interpreted in two senses: first of all, there is a group selection, where selection acts on the members of a group. This is called multi-level selection 1 (MLS1). In another sense, the group becomes the true objective of selection: in this case, selection favours group properties. This is called multi-level selection 2 (MLS2). To sum up, selection may act both within and among collectives. Cancer is both an example of MLS1 and MLS2. Firstly, selection acts among cancer cells. Moreover, selection also acts among tumour masses: cancer cells become integrated and form a group, which is the new target of natural selection. This could explain, to some extent, also the phenomenon of metastasis: "Groups are more "fit" if and only if they propagate more groups." (Lean-Plutynski 2015). In fact, cancer cells which are more successful form metastasis. In this case, fitness is assigned to the group (the tumour) as a whole: this is a true case of group selection. In fact, there is a strong cooperation among cancer cells, and variations of the traits of single cells influence the fitness of the whole tumour. This is why cancer can be told to be a product of natural selection, acting at more then one level.

At the same time, cancer is a by-product of natural selection acting at multiple levels: it uses micro-environment to grow and proliferate successfully. Therefore, for example, cancer cells acquire their phenotype thanks to the signals they both send and receive from the surrounding tissue. Cancer coopts signalling that is usually needed for the organisation of the upper level. A tumour should thus be described as a pathology which involve the disruption of hierarchical organization of metazoan.

To conclude, it will be showed how, thanks to its analysis from a multi-level perspective, cancer is increasingly considered as a pathology linked to the tissue organization more then to the cell regulation. In effect, cancer was originally considered to be a deregulation of the normal growing program of the cell. The default state of a cell was thought to be quiescence: therefore, a cell that replicates too much becomes cancerous. However, this is a simplistic view: the Darwinian explanation of cancer has given a big contribution in understanding that cancer is not only a cellular pathology. Nowadays, cancer is seen as derived from a deregulation of the connections between the tissue and the cell. This is why a cancer cell cannot be studied without any reference to its context. However, a big work still has to be done to describe the cancer cell within a larger context. In conclusion, cancer shows that biological individuality cannot be understood without referring to multiple levels simultaneously.

Foad Dizadji-Bahmani. What is the value of mathematical rigor in physics? Physics is couched in mathematics, but the use of mathematics in physics is not always uncontroversial. In particular, there are several well-known cases of physical theories (or parts thereof) that were not, at least from a contemporary mathematical perspective, mathematically rigorous. These include Newton's use of calculus in his mechanics, and the Dirac delta function in quantum mechanics. Weierstrass' work, following on from Cauchy, is usually credited as making rigorous the former. (cf. Grattan Guinness 1970) The latter was made rigorous by Schwartz in terms of distribution theory.

One thing which is uncontroversial, I think, is that mathematical rigor is, ceteris paraibus, a good thing when it comes to physical theories. It is surprising therefore that the question of what exactly is the value of mathematical rigor in physics has not been posed, let alone answered. In this paper I pose this question and try to make some headway in answering it.

Before doing so, some important caveats are needed: first, I am not assuming that mathematical rigor is necessary for a tenable physical theory (cf. Davey 2003) but just that it is a good thing insofar as one can have it. And the ceteris paribus clause is important, for it may be that rigor comes at the expense of other good making features of a theory. Second, the notion of mathematical rigor itself is unclear. The standards of mathematical rigor have shifted and rigor is not categorical. (cf. Kitcher 1981) I assume that, in any case, mathematicians or philosophers of mathematics will "know it when they see it" and I defer to their judgments about it. That is, I do not here question that, for example, Weierstrass and Schwartz did make those parts of physics more mathematically rigorous. Third, I am not assuming that there is only one kind of value in mathematical rigor for physics. There may be several distinct values; 'the' value of rigor would just be their conjunction.

In the paper I consider four kinds of value of mathematical rigor for physics: intrinsic, logical, epistemological, and representational.

Logical: the value of mathematical rigor is simply that it ensures that one avoids a logical inconsistent physical theory. It is not clear that mathematical rigor does entail logical consistency, but even if so, this defers the question. What then is the value of logical consistency? Anything follows from a contradiction of course, but it strikes me as implausible that so much intellectual labor is expended to guard against such a worry.

Epistemological: mathematical rigor bears fruit! That is, making a physical theory (more) mathematically rigorous often brings about new physical insights leading to the advancement of physics itself. Whilst there are some examples of this, increased rigor can and does cut the other way too. (Excellent discussions in Davey (2003) and Jaffe and Quinn (1993).) For this reason I argue that this does not the value of mathematical rigor.

Representational: mathematical rigor allows us to represent, through our physical theory, the actual world. The idea here is that a mathe- matically unrigorous physical theory is failing to represent a possible world or is failing to differentiate between different (sets of) possi- ble worlds which may be actual. I shall argue that this is the value of mathematical rigor and that this fits with the practice of physics, and explains why the extent to which rigor matters depends on the kind of physical theory one is considering (specifically whether it is a fundamental or phenomenological theory).

In the final part of the paper, I present an argument in favor of scientific realism: I argue that scientific realism best accounts for the value of mathematical rigor, as per the above.

Juliusz Doboszewski. Laplacian indeterminism in spacetimes with boundary There are general relativistic spacetimes in which an interesting form of indeterminism is present due to existence of (asymptotic) boundary at

timelike infinity. This indeterminism resembles the space invaders scenario, namely new information can always arrive from the boundary and influence events in the interior of spacetime. More formally, in such spacetimes for any choice of achronal submanifold S there exists an inextendible timelike curve which is not in the domain of dependence of S. Such a spacetime is not globally hyperbolic, and cannot be represented as S x R, where S is a Cauchy hypersurface (Hawking and Ellis [1973]). However, it can be foliated by global time function. Existence of such a foliation provides a friendly environment to the Laplacian analysis of determinism (the idea that given the initial state of the physical system at a moment of time, all developments of the initial state in accordance with the laws of nature are isomorphic (Earman [2007])), since the global time function is a suitable candidate for the moments of time.
Could a spacetime with (asymptotic) boundary of the kind we described be Laplacian deterministic? We can fix the boundary conditions which control hypothetical new information coming from infinity. Instead of initial value problem we should then consider initial and boundary value problem. And if that problem has unique solution, it would suggest that Laplacian determinism holds in such a spacetime. Indeed, there exists a choice of boundary condition under which the initial and boundary value problem does have a unique solution (Friedrich [2014]).
But there are two complications. First, the choice of boundary conditions in such spacetimes is highly non-unique (Friedrich [2014]). Second, the physical interpretation of the boundary conditions is unclear. Fixing the boundary condition or asymptotic behaviour of the spacetime metric amounts to an idealization (e.g. perfectly rigid boundary of the box of gas), or to picking a convenient description of some physical feature relevant to the situation at hand (e.g. asymptotic flatness represents negligible gravitational effects on the system from external sources) (Wilson [1990]). But the boundary in such spacetimes is unlike the boundary condition in these two cases. This boundary is not a location in spacetime (in particular there is no curve with bounded total acceleration which reaches the boundary). As a consequence, the boundary does not represent a physical feature of the system in the same sense as e.g. asymptotic flatness does. And because it is determined by a spacetime metric, it does represent an idealization either. Rula and Sarbach [2011] even state that it is merely a calculational device as opposed to boundaries which do have physical significance.
What are the consequences for the Laplacian analysis of determinism in such spacetimes? One of the following is the case:
(a) Laplacian determinism fails in such spacetimes, because there is no physical interpretation for the boundary conditions;
(b) Laplacian determinism holds, since one does not need to bother with providing a physical interpretation for the boundary condition;
(c) Laplacian determinism holds iff there is an interpretation of a boundary condition (preferably one which picks up exactly one boundary condition), under the choice of which the uniqueness and existence can be proven.
I will argue that (b) is not the case, because ignoring the demand for interpretation does not rule out the possibility of new information arriving, which gives the spacetime its indeterministic flavour. So either Laplacian determinism fails in such spacetimes, or a physical interpretation of the boundary condition needs to be provided. This constitutes an interesting challenge for Laplacian analysis of determinism in classical general relativity.

John Dougherty. Equality and separability in gauge theories Debates over the interpretation of gauge theories often take the main choice to be between a "fiber bundle interpretation" and a "holonomy interpretation". The former is what you get when you write down a gauge theory in terms of fields (i.e., sections of fiber bundles) and say "the world is like that". The latter is what you get when you do the analogous thing with loops in spacetime and their properties (i.e., holonomies). Healey (2007) has articulated and defended a particular holonomy interpretation of gauge theories, similar to the interpretation Belot (1998) gives in the restricted case of electromagnetism. Importantly, Healey's interpretation covers all gauge theories and extends to the quantum realm. Our interest in gauge theories stems from their role in the Standard Model of particle physics, the keystone of our theory of the small and fast. And in this role, gauge theories are quantized. So a satisfactory interpretation of gauge theories must cover the quantized case as well as the classical one.
The most striking metaphysical feature of Healey's interpretation is its nonseparability. The electromagnetic (for example) state of the world does not supervene on the states of the world's parts. Entanglement phenomena make us expect this in quantum cases, but Healey's interpretation is nonseparable even at the classical level. This is a serious metaphysical consequence. It is incompatible with the doctrine of Humean supervenience and therefore, it seems, any Humean understanding of gauge theories—though as Myrvold (2011) shows, the situation is not as bad as it seems.
But the choice between fiber bundles and holonomies is not the only choice, and focus on this distinction has obscured others. A more important choice is between a "localized gauge properties interpretation" and a "non-localized gauge properties interpretation", to use Healey's (2007, 55) terminology. The former option is often identified with the fiber bundle option, and the latter with the holonomy option, and this makes it easy to think that any holonomy interpretation is non-localized in this sense. The identification also suggests that one can dissolve the debate by showing that fiber bundles and holonomies are appropriately equivalent. Not so. Healey's interpretation of classical gauge theories is a non-localized interpretation, as evinced by its nonseparability. However, his quantum holonomy interpretation is obtained from a localized, separable classical interpretation. Every extant quantization procedure—whether it involves connections or holonomies—makes crucial use of a localized classical theory. In the holonomy representation, the localization corresponds to the inclusion of a basepoint. In the fiber bundle representation, localization corresponds to gauge freedom. Given the primacy of the quantum in our motivations, this means that we should reject the non-localized holonomy interpretation and nonseparability along with it.
Abstractly, the difference between localized and non-localized theories can be captured by whether their mathematical representation satisfies the sheaf condition, a property found in algebraic geometry. More precisely, it satisfies a generalization of this condition found in higher category theory. This naturally leads to an interpretation in which gauge-related potentials are the same state of affairs, but that avoids the problems usually associated with the simple quotient by the gauge relation. This interpretation is simply expressed in homotopy type theory, a new foundational program in mathematics.
We undertake conceptual investigation into gauge theories to understand their contributions to our best physical theories. Because gauge theories make this contribution after quantization, the primary constraint on an interpretation of classical gauge theories is that it be quantizable. This constraint has not been met by any non-localized interpretation, including Healey's. This suggests that one notion of separability is important for quantization, even though another kind is violated in Bell-type experiments. The focus on differences between fiber bundle and holonomy interpretations has obscured this distinction, to our detriment.

Douglas Earl. On the Concept of Superposition in Quantum Field Theories The concept of superposition is used in two ways in QFT: First, to analyse free-field states in terms of particles (or quanta) and their states where it may be assumed that quanta are sufficiently separated to neglect their interactions; Secondly, to analyse and explain interaction processes. Interactions are often understood to occur as a superposition of exchanges of 'virtual' quanta as represented by Feynman diagrams (Brown & Harré (eds), 1988; Falkenburg, 2007). Through an analysis of 'superposition' from its origins in classical mechanics through quantum mechanics to QFT I show that it is correctly applied in free-field cases, but misapplied in the analysis of interactions. This has lead to metaphysical confusions regarding the nature of interactions.
To get a handle on the notion of superposition in QFT, I will compare it with uses of 'superposition' elsewhere in physics. First, I consider superposition in classical mechanics. Here it is associated with systems described by linear differential equations, especially those that behave in wave-like ways. The vibrations of a stretched string as modeled by the one-dimensional wave equation offers a paradigmatic example. The superposition principle states that any linear combination of solutions to a linear differential equation modeling a system is also a solution. In many cases such equations admit a set of 'basic' solutions, the normal modes. Any solution can be written as a superposition of normal modes. This has close affinities with Fourier analysis and Fourier series solutions of differential equations in which the Fourier modes are the normal modes. So for instance for the stretched string the normal modes are the Fourier modes, which are well-defined as the string's harmonics.
The key points to be developed here are that: (i) Superposition, and normal mode decomposition, are concepts associated with linear differential equations: they are inapplicable in non-linear cases (Dunn, 2013); (ii) The identity and persistence conditions of system behaviours modelled via certain linear differential equations are given via superpositions of the Fourier modes; (iii) Each Fourier mode individually is a solution to the differential equation, representing a possible system behaviour. In this sense the modes individually have physical significance. So for the stretched string the modes (i.e., the harmonics) are individually possible motions of the string, any motion of the string can be written uniquely as a superposition of these modes, and the modes persist in the behaviour of the string.
It might be that in a robust metaphysical sense any physical state of such a modelled system is a composition of modes, but I shall not develop or use this strong claim. Rather, I explore the weaker claim that a realist commitment to a term in a mathematical representation of the behaviour of a system requires the term individually to have physical significance as just outlined. This enables the demarcation of my approach to superposition from a realist interpretation of epicyclical analysis of planetary motions, or arbitrary mathematical 'decomposition' of a continuous mass as a sum of masses.
Moreover, this move enables a philosophically important contrast to be drawn between modal or Fourier analysis and power-series or iterative solution techniques to differential equations. Returning to the stretched string as an example, the one-dimensional wave equation has a Fourier series solution as a superposition of modes. Each series term or mode represents a possible string behaviour, and so has physical significance. This is contrasted with both iterative and power-series solutions. Both methods result in a power series giving the exact solution in the limit. The key point is that unlike the Fourier case, no individual term of the power series is a solution to the equation. The individual power-series terms lack physical significance, and the power-series solution cannot be described as a superposition of terms.
Secondly, brief consideration is given to the concept of superposition in quantum mechanics. The concept was central in Dirac's exposition (1958), but absent in von Neumann's (1932). Mathematical treatments of quantum theory can be given either with or without reference to superposition, suggesting that it pertains to the physical interpretation and significance of states or wavefunctions. Superposition is used in conjunction with descriptions of eigenstates and solutions to Schrödinger's equation. Schrödinger's equation is a linear differential equation, and the concept of superposition is properly used in continuity with its classical usage. The eigenstate solutions are (generalised) Fourier modes. What differs in the quantum case and remains open to development (on the Copenhagen interpretation) is the probabilistic interpretation of terms in a superposition and their dependence on the observable at hand. This will not be explored here.
These clarifications are brought to bear on QFT. Initial and final states are modelled by free-field, linear differential equations. Such states are capable of interpretation as superpositions of Fourier modes identified as field quanta (e.g. electrons). This gives identity and persistence conditions of field states in terms of 'particles'. So one may identify and count electrons 'as if' they existed as individuals. This is in continuity with classical cases. The interaction process is, however, modelled by a non-linear equation, so one should expect the concepts of modal decomposition and superposition to be inapplicable. This is explicitly stated in the context of non-linear differential equations modelling systems in mechanical and electrical engineering which are often solved with Volterra series methods (Dunn, 2013). Indeed, in QFT Dyson's expansion, which is an iterative, asymptotic series approximating the solution of the key differential equation modelling interaction, is formally similar to a Volterra series. Individual terms in Dyson's series are associated with Feynman diagrams and exchanges of virtual quanta. These are best compared with individual terms in the power-series solution to the wave equation rather than with the Fourier mode decomposition. Individually they lack physical significance since no individual term in Dyson's expansion is a solution to the equation. One cannot identify quanta as modes as in the free-field case. It is meaningless to talk of superpositions of virtual quanta exchange processes, so a different explanation or metaphysical picture of interaction is required, whilst the empirical adequacy of the theory can be explained.

Matthias Egg. Real Patterns without Underlying Stuff A central issue in the ontological debate on quantum mechanics is the question whether the high-dimensional space on which the quantum mechanical wave function is defined (the so-called configuration space) represents the fundamental space in which physical reality unfolds. Whoever wants to give an affirmative answer to that question faces the challenge of recovering our experience of a three-dimensional world from a high-dimensional underlying ontology. In my talk, I will discuss (and sketch a solution to) two problems that affect the two most popular attempts to do this (proposed by David Wallace and David Albert, respectively; see, e.g., Wallace 2012, chapter 2 and Albert 2015, chapter 6).

Despite considerable differences in focus and context, Wallace's and Albert's proposals share the basic idea that the three-dimensional structures with which we are familiar are to be viewed as certain dynamical patterns in the wave function of the universe. A first problem with this view, as pointed out by Alyssa Ney (2015, p. 3118), is that the notion of a pattern seems to presuppose the existence of some underlying stuff, something of which there is a pattern. Indeed, Wallace borrows the relevant concept of a pattern from Daniel Dennett (1991), and all the examples Wallace and Dennett use to illustrate this idea deal with patterns as spatiotemporal arrangements of some kind of matter. But what is required if one starts with nothing but the wave function is a story of how we get matter (situated in ordinary space and time) in the first place, so its existence must not be presupposed in the concept of a pattern.

In response, I first remark that there is no conceptual obstacle to applying the notion of a pattern beyond the realm of spatiotemporally located entities. For example, we readily understand what it means for there to be a pattern in a sequence of natural numbers. What is needed, then, in order to overcome Ney's criticism, is a sufficiently precise way to explicate this extended application of the pattern concept beyond the spatiotemporal realm. I will argue that the technically precise notion of a real pattern developed by James Ladyman and Don Ross (2007, chapter 4) is fit for that task, although some interpretive work needs to be done to deal with an apparent tension between their information-theoretic account of patternhood and the strongly realistic interpretation of the wave function that seems to be presupposed in Wallace's and Albert's accounts. I will therefore try to demonstrate that, although Ladyman and Ross do not admit any fundamental substance in their ontology, arguing instead that "it's real patterns all the way down" (2007, p. 228), their account of real patterns is ontologically flexible enough to be combined with wave function realism in the sense of Wallace and Albert.

In his discussion of Albert's proposal, Peter Lewis (2013, p. 116) emphasizes a second problem for the appeal to patterns in the wave function as an explanation of three-dimensionality. He complains that the patterns on which Albert relies for identifying three-dimensional structures appear only under one arbitrary choice of coordinates, and such patterns are generally regarded as artefacts of that choice rather than facts about the world. This would render the pattern story rather unattractive, at least to those who want to maintain realism about ordinary physical objects.

Like the first problem, this second one can also be traced back to Dennett's (1991) seminal paper, which, despite its title ("Real patterns") has often been interpreted as advocating an instrumentalist (rather than realist) stance towards the patterns. Against this, Ladyman and Ross (2007, chapter 4) argue for a realistic reading of Dennett, and I will apply this realism in order to counter Lewis's critique. By means of a simple mathematical example, I will illustrate how the coordinate-dependence to which Lewis alerts us is compatible with realism. The point is that whether or not a pattern appears to us may depend on the choice of coordinates, but whether or not it is there does not.

The upshot of this discussion is that the alleged obstacles to identifying the objects of our experience with patterns in the wave function can be overcome.

Joshua Eisenthal. Ontological vs. Logical Foundations in Hertz's Mechanics On new year's day of 1894, Hertz died just 36 years old. He had been heralded as one of the leading minds of his generation, and many had looked to him in expectation of major new discoveries. Ernst Mach paid tribute in his memorial address: 'Heinrich Hertz... one of the leaders of our discipline, the pride and hope of our country, has been committed to the grave.' Hertz had dedicated the last few years of his life to a grand project in foundations of physics, culminating in the posthumous publication of "Principles of Mechanics". When, with much anticipation, the book finally appeared, it was received with much praise. But even as it was admired for its elegance and scope, Hertz's contemporaries could not find in it the kinds of advances and insights that they had hoped for. Indeed, there was a general sense of confusion regarding what it was that Hertz had taken himself to have achieved. Hertz himself, of course, could not help. As Boltzmann put the matter, 'at the same moment his lips became for ever sealed to the thousand requests for clarification that are certainly not on the tip of my tongue alone'.

One clear incentive to search for a more satisfactory interpretation of Principles stems from the fact that it has had a significant influence on subsequent philosophy. An important but poorly understood example is the influence on Wittgenstein, who had a deep and lifelong appreciation of Hertz's book. Wittgenstein referred to it twice in the Tractatus, and even considered using a quotation as the motto for the Philosophical Investigations. He also quoted from Principles both times that he gave a programmatic address as a Cambridge professor. The link with the Tractatus is particularly striking. Both texts seem to present an ontological picture founded on fundamental simples. For Hertz, 'material particles' make up mechanical 'systems', and for Wittgenstein, simple 'objects' make up 'states of affairs'. However, both authors appear remarkably unconcerned to present any concrete examples of such entities. Thus a problem in interpreting either text is working out what such entities might be, on the one hand, and understanding the seemingly cavalier attitude of the authors in this regard, on the other.

Within Tractatus scholarship, this is a particular problem faced by "ontologically-oriented" interpretations which seek to understand Tractarian objects as ontological simples, such as sense-data or point particles. Unsurprisingly it is in this context that appeals to Hertz's influence have sometimes been made, with the proposal that the material particles in Principles were the simple objects that Wittgenstein had in mind. However, an increasing number of Wittgenstein scholars have come to regard ontological interpretations of Tractarian objects as fundamentally misguided. A contrasting "logically-oriented" approach emphasises the process of logical segmentation that is inherent in logical analysis, as captured in Frege's method of arriving at his object-concept distinction via the analysis of judgment: 'instead of putting a judgment together out of an individual as subject and already previously formed concept as predicate, we do the opposite and arrive at a concept by splitting up the content of possible judgment'. Logically-oriented commentators who take this to inform Frege's context principle – 'never to ask for the meaning of a word in isolation, but only in the context of a proposition' – draw attention to Tractarian versions of the principle such as that at 3.3: 'Only propositions have sense; only in the nexus of a proposition does a name have meaning.'

To return to Hertz, such logically-oriented commentators have generally not regarded the details of Principles as standing to contribute to our understanding of Tractarian objects. Indeed, it might seem prima facie implausible that Hertz's treatment of physics would be analogous, in this particular way, with Frege's and Wittgenstein's treatment of language. However, a careful reading shows that Principles directly invites a logically-oriented interpretation. For example, in Hertz's framework it is complex systems that are given priority over simple points: 'in reality, the material particle is simply an abstraction, whereas the material system is presented directly to us... it is only by processes of reasoning that we deduce conclusions as to possible experiences with single points' (Principles, p.31). Despite the fact that Hertz emphasised that his purpose was to present a logically perspicuous reformulation of the known content of classical mechanics, it is clear that many commentators missed his intent. A major source of confusion is one of the key characteristics of Hertz's formulation – the fact that he makes the notion of distance forces derivative on the notion of 'connections'. The simplest example of a connection is a rigid constraint connecting two material points, but at its heart Hertz's notion pertains to analytical representations of a system, i.e. permissible changes of coordinates within a mathematical description. Indeed, it would be disastrous to regard Hertz's connections as actually linking physical objects, for this would make his formulation hopelessly implausible. Yet some such "ontological" interpretation is what seems to have confused Helmholtz, Mach, Boltzmann and others. FitzGerald provides perhaps the most vivid example of an ontologically-oriented reading of Principles in remarks likes the following: '[Hertz] does not seem to investigate anywhere the question as to the danger of his rigid connections becoming tangled.'

In this paper I will outline a logically-oriented interpretation of Principles, and argue that such an approach is the key to making sense of Hertz's project. Furthermore, I will suggest that it is only on this basis that we can appreciate Hertz's influence on Wittgenstein. I will begin by accounting for the reception of Principles by Hertz's contemporaries, exploring their widely shared ontologically-oriented interpretation of the text against the background of nineteenth century efforts to model the luminiferous ether. I will then highlight the appeal of a logically-oriented interpretation of Principles by tracing the development of Hertz's thought concerning scientific representation, and outline what such an interpretation would come down to.

Peter Evans. Fluid mechanical models and causally symmetric approaches to quantum mechanics A recent series of experiments from a team in Paris have demonstrated that an oil droplet bouncing on a vibrating fluid surface displays behaviour that is typically considered to be quantum behaviour. This behaviour includes single and double slit diffraction and interference (Couder and Fort, 2006) and quantised orbits of bound state pairs (Fort et al., 2010), as well as phenomena that look analogous to quantum tunnelling (Eddi et al., 2009), Schrödinger evolution of probabilities (Couder and Fort, 2012) and Zeeman splitting (Eddi et al., 2012). There are two key features of these experiments. The first is that each droplet is a local source of a standing Faraday wave that is sustained by the externally driven, vertically vibrating fluid (Protière et al., 2005). Given the right conditions, this Faraday wave then propels the droplet to "walk" horizontally across the surface, coupling the motion of the droplet ("the particle") to the vertical displacement of the fluid surface ("the wave") (Protière et al., 2006). The second key feature is that, due to the external impetus applied to vibrate the fluid, the standing Faraday waves created from each bounce of the droplet are sustained (again, in the right experimental regime) for very many bounces of the droplet. The vertical displacement of the fluid surface at any point is thus the linear combination of very many distinct Faraday waves, and this provides "the wave" with a path memory of where "the particle" has recently been (Eddi et al., 2011).

The suggestion has been made that this fluid mechanical system provides a single particle classical model of the pilot wave mechanism of the De Broglie-Bohm theory of quantum mechanics. The claim is that the wave on the fluid surface acts as a pilot wave that guides the behaviour of the oil droplet in analogy to the way in which the De Broglie-Bohm pilot wave guides the behaviour of the particle configuration (in this case, a single quantum particle in 3-space). It is interesting to inquire, though, into the relation between the explicit nonlocality of the De Broglie-Bohm pilot wave, and the patent locality of the wave on the fluid surface (both driving apparent nonlocal behaviour in the particle).

A better understanding of the relation between these local and nonlocal elements of the fluid mechanical model and quantum mechanical theory can be gained by considering a further feature of the Paris experiments that has as yet flown somewhat under the radar (although Vervoort (2015) gestures towards it). Each time the droplet bounces on the fluid surface a distinct damped travelling capillary wave is emitted that travels at a velocity typically about 10 times the walking velocity of the particle (Protière et al., 2006, p. 95). Given a boundary feature---such as a wall, a slit, a submerged barrier or a new particle with its own associated wave---within the damping length of the capillary wave, the standing Faraday wave that results from the superposition of successive capillary waves will contain a reflected component encoding information about this boundary. Since the particle is coupled to this wave, this accounts for how the motion of the particle can be influenced by spatially remote boundaries (so long as they are within the damping length), and this produces the apparent nonlocal behaviour of the particle from a local mechanism. (This is manifest in, for instance, the single slit diffraction experiments when the particle begins to deviate from its straight path well before it reaches the aperture, and this deviation at least partly accounts for the subsequent diffraction.)

Significantly, some of these nonlocal boundaries (including the single slit) lie in the 'future' path of the walking particle. This enables the fluid mechanical model to gain a sort of 'local advantage' over De Broglie-Bohm theory: by seeking out relevant 'future' boundaries, the fluid surface wave can locally encode information about distant spatial boundaries, all made possible by the relatively slow velocity of the particles compared to the probing capillary waves it creates on the fluid surface. But this is then immediately recognisable as the feature that sets apart causally symmetric approaches to quantum mechanics, and there are two in particular that this model resembles most closely: the transactional interpretation of quantum mechanics (Cramer, 1986), and the causally symmetric Bohm model (Sutherland, 2008).

This talk explores to what extent the fluid mechanical model of the Paris experiments can be interpreted as providing a causally symmetric model of quantum mechanics. I contend that the Paris experiments provide a direct model for neither the transactional interpretation nor the causally symmetric Bohm model, but rather for a Frankenstein amalgamation of the two.

Matt Farr. Causation and Time Reversal 1. Overview.

The question (What would the world be like if run backwards in time?) is ambiguous: does a 'backwards-in-time' world involve an inversion of cause and effect? In the case of time-reversal invariant theories, it is common to understand time reversal 'causally', holding that whatever can happen forwards in time can happen backwards in time. This causal interpretation of time reversal is problematic as it appears to be incompatible with the asymmetry of cause and effect, and hence has been widely taken to motivate eliminativism about causation. I argue that such worries are misplaced on two grounds. First, a 'causal' interpretation of time reversal is poorly-motivated. I argue that time reversal should be understood 'non-causally', such that a world 'run backwards' is not a genuine possibility — pairs of worlds that are the time-reverse of each other contain the same causal relations. Second, I show that even on a causal interpretation of time reversal, a time reversal symmetric theory is compatible with causation.

2. The Directionality Argument.

Bertrand Russell famously took time symmetric features of the law of gravitation (taken by Russell as an exemplar of physical laws) as incompatible with the asymmetry and time asymmetry of causation. This has recently been generalised into the 'Directionality Argument' which holds that time reversal invariance is incompatible with causation (cf. Ney 2009, Norton 2009, Frisch 2012).

3. What does time reversal reverse?

To assess the compatibility of time-reversal invariance and causation, I focus on whether time reversal should be taken to invert causal relations, giving us two options:

- *Causal time reversal.* Time reversal involves inverting causal relations, taking causes to effects and vice versa.
- *Non-causal time reversal.* Time reversal does not invert causal relations; the distinction between cause and effect remains invariant under time reversal.

I show that these sit best with distinct ontologies of temporal relations. First, causal time reversal implies a 'B-theory' of time - an ontology of time-directed relations, such that two worlds may differ solely with respect to the direction of time. Second, non-causal time reversal implies a 'C-theory' of time - an ontology of undirected temporal relations, such that no two worlds may differ solely with respect to the direction of time. On the C theory, inverting every 'earlier than' relation gets back the same possible world. Only on the C theory does time reversal amount to a redescription of a single causal structure.

4. Does time reversal symmetry eliminate causation?

Using examples of time symmetric and time asymmetric processes, I argue that the C theory provides superior accounts of causal relations and time reversal. I'll here consider the time-asymmetric example. Imagine three pictures side-by-side depicting the time asymmetric collision of two snooker balls of equal mass on a frictional snooker table. The leftmost picture, L, represents both the cue and object (red) ball at rest 30cm apart with a cue primed to strike the cue ball towards the red ball. The centre picture, C, represents the cue ball with velocity v at the instant of collision with the red ball. The rightmost picture, R, depicts both balls at rest, with the red ball displaced one metre to the right of its initial position and the cue ball at rest at the point of impact. In the conventional L–R description, the cue strikes the cue ball, setting it in motion, and the cue ball collides with the red ball, transferring most of its momentum to the red ball and subsequently being at rest relative to the table due to spin. The red ball then loses momentum due to the frictional force of the table until at rest, as in R. The L–R description contains a number of causal terms, implying: the cue movement causes the cue ball's movement; the cue ball's movement causes the red ball's movement; the baize causes the red ball to lose momentum. In the unconventional R–L description, an anomalous series of causal processes is implied. Firstly, heat in the baize together with incoming air molecules conspire to set the red ball in motion. Secondly, the red ball's motion in synchrony with inverse, concentrating soundwaves jointly impart a gain in momentum in the collision of the red ball into the cue ball. Finally, the cue ball's momentum is absorbed in a collision with the cue.

5. Answers.

On the C theory both the forwards and backwards descriptions pick out the same causal structure. On the B theory, the two descriptions pick out distinct causal structures. I argue that: (1) the C theory's non-causal account of time reversal is clearly preferable on pragmatic grounds given consideration of time asymmetric causal processes; (2) on neither the causal nor non-causal readings of time reversal is time reversal invariance incompatible with causation.

First, as a candidate causal process, R–L is unsatisfactory. Two issues stand out: (a) R-L implies violations of the Causal Markov Condition; (b) the snooker player loses her agential control over the balls' motion. Only on the B theory is R-L a distinct candidate causal process from L–R.

Second, for causation and time reversal symmetry to be incompatible is for a time reversal invariant theory to entail, for some events x and y, that if x is a cause of y, then (by time reversal symmetry) y is also a cause of x (and vice versa). By holding time-reversal-related models to represent distinct worlds, the B theory avoids incompatibility since the distinct causal relations expressed by time-reversal-related models hold between distinct sets of events in different possible worlds. On the C theory, incompatibility is avoided on the grounds that time reversal does not invert causal relations. Moreover, the non-causal account of time reversal entails this conclusion regardless of whether the relevant physics is time reversal invariant.

Laura Felline. It's a matter of Principle. Explanation in Axiomatic Reconstructions of Quantum Theory in terms of Information-theoretic Principles. Within the current mainstream research in the foundations of QT much attention has been turned to Quantum Information Theory (QIT), and in particular to those reconstructions that focus on information-theoretic principles (Grinbaum, 2007). It is therefore becoming increasingly important for philosophy of science to deal with the explanatory gain that is implied in the switch from the traditional interpretation program to the program of Axiomatic Reconstruction of QT in terms of Information-theoretic principles (ARQIT, henceforth).

In this paper I explore ways in which ARQITs could contribute to explaining and understanding quantum phenomena, as well as studying their explanatory limitations.

As a concrete case study, I analyse how ARQITs can account and explain quantum non-locality, defined minimally as non-local quantum correlations in the sense of Bell's theorem. To make the analysis more concrete, I take as an illustrating example a specific ARQIT, i.e. Clifton, Bub and Halvorson (CBH) characterizing theorem (Clifton et al. 2003). However, the analysis here proposed can be extended, mutatis mutandis, to other reconstructions.

The CBH theorem presupposes a mathematical background called C*-algebra which, according to the authors, is neutral enough to allow a mathematically abstract characterization of a physical theory. Within such a framework, CBH formulate a theorem characterizing QT in terms of three principles about impossibilities of information transfer:

The impossibility of superluminal information transfer between two physical systems by performing measurements on one of them.

The impossibility of perfectly broadcasting the information contained in an unknown physical state.

The impossibility of unconditionally secure bit commitment.

More specifically, the explanation of non-locality follows three steps:

1. The 'no superluminal information transfer via measurement' implies the commutativity of distinct algebras. Commutativity of distinct algebras is meant to represent no-signalling. A theory violating this principle would display strong non-locality and superluminal signaling.
2. The 'no broadcasting' principle implies the non-commutativity of individual algebras. A theory violating this principle is therefore a classical theory with commutative individual algebras.
3. CBH show that and how quantum non-locality follows from these two joints algebraic features: if A and B are two sub non-commutative mutually commuting algebras, there are nonlocal entangled states on the C*-algebra A VB they generate.

In developing the details of my account, I take inspiration from the account of explanatory proof in mathematics formulated by Mark Steiner (1978). Although Steiner's account was developed as an account of explanations in pure mathematic, its central idea that "to explain the behavior of an entity, one deduces the behavior from the essence or nature of the entity" (p.143), captures what (I submit) is the explanatory content of ARQITs. Since talking of the 'essence' or 'nature' of a mathematical (as much as of a physical!) system is quite problematic, Steiner appeals to the concept of characterizing property, i.e. a "property unique to a given entity or structure within a family or domain of such entities or structures." (Steiner 1978, p.143) I will argue that the definition of characterizing property applies to the principles of ARQIT, whose function is to isolate QT against a family of theories. The CBH theorem, for instance, isolates QT against the family of all mainline physical theories, which CBH take to be all theories representable with a C*-algebra.

An essential part of the explanation of ARQIT consists in the derivation of the explanandum, which, given the set of QT principles (or axioms), is shown to be a theorem of QT. For instance, the explanation of non-local entanglement provided by the CBH reconstruction consists partly in making explicit how the existence of entangled states follows from no superluminal signals and no broadcasting.

Derivation, however, is not sufficient. For the comprehensiveness of the explanation it is crucial to show how, by changing the characteristic property, the theorem/explanandum changes in response. As in Steiner's account, here "explanation is not simply a relation between a proof and a theorem; rather, a relation between an array of proofs and an array of theorems, where the proofs are obtained from one another by the 'deformation' prescribed above." (p.144) In CBH, for instance, if no broadcasting is dropped, then one has a classical phase space theory, while if the no-superluminal signals is dropped, one has a theory where distinct and distant physical systems are not kinematically independent, i.e. a strongly non-local theory.

More concretely, under the account I am proposing, ARQIT's explanations provide a specific kind of 'what-if-things-would-have-been-different' knowledge, where the counterfactual claims are produced by variating the theory's principles and by mathematically deriving what consequences follow from such a deformation.

This suggests that this variety of explanation naturally fits in recent accounts of scientific explanation that converge in attributing a central role to the counterfactual dependence between explanans and explanandum (Morrison 1999, Bokulich 2009, Reutlinger 2012).

Finally, given the features of the described explanation, I will argue that this one provided by ARQIT is the explanation of some aspects of non-locality. I will counter the claim that with such an explanation ARQIT makes a traditional interpretation of QT explanatory irrelevant, and the latter should be therefore straightforwardly ruled out (Clifton et al., 2001).

Enno Fischer. Tacit Knowledge in Science It is widely believed that experimental results have to fulfil certain criteria of objectivity and universality in order to count as evidence for a scientific theory. In particular, it is held that scientists who try to establish a theory have to supplement their experimental evidence by sufficient details as to enable any other sufficiently trained scientist to repeat the experiment. Hence, replicability is considered a necessary condition for the objectivity of experimental evidence.

However, on the example of reproducing TEA-lasers Collins showed, that not even specialists in laser-building were able to reproduce this new kind of laser only by reading published instructions. His conclusion was that replication is a matter of experimental skill which includes tacit knowledge (Collins 1985).

With 'tacit knowledge' Collins took up a concept which was previously coined by Michael Polanyi. But from Collins' early work on tacit knowledge it is obvious that his notion of tacit knowledge is not congruent with Polanyi's approach. In fact, Collins' early use of this term is ambiguous. This is a problem that he tried to resolve in a more recent work (2010) which gives a more elaborate account of tacit knowledge. However, the proposed disambiguation of 'tacit knowledge' is not very useful in the context of Collins' earlier findings and still does not take Polanyi's approach into account.

In my talk I want to advance a new taxonomy of tacit knowledge which can be summarised as follows:

(WTK) Weak tacit knowledge is knowledge which is as a matter of contingency not articulated at the moment. For example when a person explains a game to another person she may forget to mention one of the rules of the game. In this case the rule remains tacit although there are in principle no obstacles to an articulation.

(ITK) Any aspect of intermediate tacit knowledge can be made explicit (like WTK). However, as soon as any aspect is made explicit it will point towards new aspects which are not yet made explicit. Consequently, the content of ITK can be approached through articulation but it cannot be made completely explicit.

(STK) Strong tacit knowledge is more precisely described as tacit knowing since it does not involve content which could be disclosed (like WTK) or approached (like ITK). On the contrary, it is a mode of being aware of particular elements of an action, of perception or of scientific thinking. Any attempt of making the particular elements explicit results in leaving this mode.

I argue that this taxonomy gives fresh insight into two issues in the philosophy of science. First, Collins' problem of replication and the experimenters' regress

(1) and, secondly, Duhemian underdetermination (2). Moreover, the taxonomy relates Collins' earlier findings to Polanyi's original notion of tacit knowledge which can be identified as STK (3).

(1) I suggest that knowledge transfer in experimental contexts is invisible and cannot be guaranteed (as Collins claims) since it involves ITK. This implies that we cannot give a comprehensive set of criteria which upon fulfilment guarantee successful replication. Therefore, we do not know whether an attempt of replication is successful unless and until the previous result has been reproduced. Thus, any attempt of replication can only be assessed on the basis of the outcome of the original experiment. This requires that the validity of the original experiment is presumed. But this is a circular criterion because the validity of an experiment depends on its replicability.

The circularity is not vicious as long as all involved scientists agree upon what counts as the correct outcome of the original experiment. However, it is vicious when the experiments deal with controversial new findings. Then the circularity brings about what is known as the experimenters' regress: the circularity can only be broken through an independent method which assesses the validity of the original experiment and its replication (calibration).

(2) The difficulties in agreeing about experiments and their outcomes which prevail in the context of replication affect also Duhemian underdetermination. When an experiment is taken as a test on a theory we do not only rely on the laws of physics which are employed in the construction of measurement devices but also on a group of facts regarding the proper conduction of the experiment.

I argue that the relevance of ITK in the context of underdetermination suggests that the experimenters' regress can be seen as consisting of subsequent stages of underdetermination which are interconnected through the introduction of calibration methods.

(3) Polanyi (1962) identifies tacit knowledge (STK in my taxonomy) through an analogy to perception. In the optical illusion of the Ames room we perceive two seemingly different sized objects on the background of a seemingly rectangular room. By a change of perspective we can easily see that the objects are, in fact, equal in size and that our previous perception was distorted by the background which is in fact trapezoidal. Polanyi concludes that whenever we are focally aware of something (the seemingly equal-sized objects), we are relying on our subsidiary awareness of a background which gives us clues (seemingly rectangular-shaped background). The perception of clues is influenced by the whole history of our previous experience (we usually inhabit rectangular-shaped rooms).

This has two consequences. First, we often overlook the unprecedented. In the context of science this means that we often are able to find regularities in our experience only where we expect them to occur. Secondly, the fact that we are surprised when we see that the objects are equal in size, indicates how difficult it is to keep track of the clues that affect our perception. In fact, it is a necessary condition that the clues are not explicit for them to have suggestive power (STK). As soon as we try to make them explicit and they do not appear in terms of their function and, thus, are disintegrated. This makes it impossible to perceive new regularities.

Samuel Fletcher. On the Alleged Incommensurability of Newtonian and Relativistic Mass One of the enduring debates about scientific change concerns the extent to which there is conceptual continuity across successive theories. The same term as used in different theories often on its face appears to have ultimately different extensions. Despite some ostensive overlap, the traditional story goes, they are embedded in a different network of terms that, holistically, grants it a different meaning. There has also been a more recent resurgence of debate regarding limiting-type relationships between theories, especially in physics, and whether these count as reductive relationships. This debate has concerned to what extent one theory can be the limit of another, and whether, if it is, this explains the limit theory. Although these two debates are not always explicitly connected, one of my goals is to show how a particular sort of positive solution to the reduction question can also contribute to understanding the extent of conceptual continuity and discontinuity between theories related by a limit. In particular, I apply some relatively new (to the philosophical literature) topological tools for understanding the limiting relationship between Newtonian and relativistic kinematics to what is perhaps the most well-known alleged example of conceptual incommensurability, that between the Newtonian and relativistic concepts of mass. My main contention is that the mass concept in the two theories of kinematics is essentially the same.

Famously, of course, both Kuhn and Feyerabend provided historical evidence that, in the mathematical framework used to formulate Newtonian and relativistic kinematics at the latter's inception in 1905, these concepts were not the same. I do not intend here to dispute their historical claims. Rather, my contention is based on a reconstruction of both theories in light of the best mathematical frameworks for describing them now, that of four-dimensional differential (affine) geometry. Thus, I do not intend to dispute here how historical actors involved in the construction, elaboration, and propagation of relativity theory. Instead, I wish to show that however the situation appeared to these actors, there is a way of describing and understanding these theories and their relationship that makes completely transparent the commonality of their concepts of mass.

One of the interesting conclusions to draw from this is that the usual understanding of incommensurability is likely too tied to the contingent and accidental features of the particular language in which a theory may be described—that is, it is too tied to the syntactic conception of theories that dominated philosophy of science in the 1960s. While there continues to be debate about the merits of the semantic view of theories, the syntactic view's successor, almost all seem to be in agreement that capturing the structure of a theory involves in large part aspects that are invariant (or at least appropriately covariant) across choice of language. Taking this into account shows that the essential differences are not so invariant. This moral is important for the reduction literature, too, for one potential objection to the claim that Newtonian kinematics is the reductive limit of relativistic kinematics is that the incommensurability of their mass concepts prevents the limit from being reductive, i.e., explanatory. Thus showing the commonality of the mass concepts is also important for understanding the explanatory relationship between the theories.

The technical portion of my argument proceeds in three phases. The first involves formulating both Newtonian and (special) relativistic kinematics in the framework of four-dimensional differential (affine) geometry, with the worldlines of particles as certain (timelike) piecewise smooth one-dimensional submanifolds. In both kinematical theories, mass is a non-negative parameter that, when associated with a worldline, specifies the degree to which the worldline departs from being a geodesics—following locally straight ("unforced") motion. The mass parameter then in both theories enters into the expression of the particle's four-momentum as a kind of normalization constant. I point out that there is a degree of convention not normally recognized in how it so enters, but that the choice of convention is essentially irrelevant when considering the details of simple particle collisions. The completion of this formulation reveals that mass plays the same functional roles in both kinematical theories; the only substantive difference lies in different spacetime structures that determine spatial distances and temporal lengths.

These different structures are nonetheless related, and in the second technical phase, I show how the Newtonian structure arises at the limit of the relativistic structure. This limit is constructed mathematically, by considering sequences of relativistic spacetimes (with various particles and observers within) that converge to Newtonian spacetimes, the sense of convergence being given by an appropriately chosen topology on the joint class of spacetimes. Because the Newtonian and relativistic spacetimes have a common conceptual interpretation, as revealed in the first phase, the topology can be easily interpreted as encoding similarity of empirical predictions. Thus a convergent sequence of relativistic spacetimes does not indicate a sequence in which the speed of light grows without bound, but rather one in which the measurements of the fixed observers can be better and better approximated by those of a certain hypothetical idealized Newtonian observer.

The third phase responds to a natural objection to the above account, namely that it has not explained the significant difference of Einstein's mass-energy relation, E=mc2. Here I build on previous work by Rindler, Lange, and Flores, as well as on the conventional elements mentioned above, to explain the significance of the most famous equation not asserting the identity of mass and energy, but either as defining energy or stating an energy content associated with mass. The analysis of classical "fission" experiments can then be made where change in mass is interpreted only as a change in effective mass, a conceptual move also available in the Newtonian framework. Lastly, I gesture towards how this analysis extends to the Newtonian and general relativistic theories of gravitation, the former in its Newton-Cartan form, where the presence of the same sort of mass can be understood as having the same sort of influence on spacetime geometry.

Simon Friederich. Merits and Limits of the Fine-tuning Argument for the Multiverse Are there other universes, some of them perhaps radically different from our own? Many physicists and philosophers are attracted to this multiverse idea because they regard it as offering a promising response to the perennial puzzle that many constants of nature appear to be fine-tuned for life: had they been even slightly different, life as we know it could not have existed.

The fine-tuning of the constants has been cited as support for the idea of divine creation, but the multiverse idea may offer an attractive non-theistic alternative: if there is a sufficiently diverse multiverse, the constants may differ between universes, and it is only to be expected that there is at least one universe where they are right for life. As the notorious (weak) anthropic principle highlights, observations can only be made where conditions allow the existence of observers. Assuming that observers are living organisms, observers can only exist where conditions are right for life, so the multiverse inhabitants—if there are any—will unavoidably find the constants apparently fine-tuned for life. The multiverse idea thus provides a candidate explanation of why we exist despite the required fine-tuning of the constants and, derivatively, by appeal to the weak anthropic principle, why the constants appear fine-tuned for life.

The multiverse idea is extremely controversial among physicists and not unanimously popular among philosophers. Some reject the fine-tuning argument for the multiverse as fundamentally flawed on grounds that, according to them, it commits a version of what Hacking calls the inverse gambler's fallacy. The inverse gambler's fallacy is committed by someone who infers from witnessing a remarkable outcome (a triple six in a triple coin toss, say) in a series of trials in a random procedure that the overall number of trials is (or has been) likely large. This inference is fallacious because the trials are probabilistically independent, which means that the existence of other trials is evidentially irrelevant for the outcome of the present one. According to White, the fine-tuning argument for the multiverse is guilty of this fallacy by "supposing that the existence of many other universes makes it more likely that this one—the only one that we have observed—will be life-permitting." (White 2000, p. 263) White's criticism, which has become known as the this universe objection, has been endorsed by other philosophers as well.

However, not all philosophers agree with the this universe objection against the fine-tuning argument for the multiverse. Darren Bradley (2009), for example, argues that the objection overlooks the unavoidable bias that results from the fact that, as the weak anthropic principle highlights, we could not possibly have found ourselves in a life-hostile universe. According to him, if we take this bias into account, it becomes clear that the fine-tuning argument for the multiverse is valid and does not commit the inverse gambler's fallacy. However, Landsman (forthcoming) has recently disputed the adequacy of the analogies appealed to by Bradley, and the debate seems to have reached deadlock.

In this contribution, I take a fresh look at the fine-tuning argument for the multiverse by considering two other problems of apparent fine-tuning, which are in many ways analogous to the problem of fine-tuned constants: the problem of our apparently fine-tuned planet and the problem of our apparently fine-tuned ancestors. Table 1 (see below) gives an overview of the three fine-tuning problems and the corresponding "many ..." responses to them. As I shall argue, it is uncontroversial that reasoning about the two other problems which parallels the fine-tuning argument for the multiverse does not commit the inverse gambler's fallacy. This allows to make a confident claim that the fine-tuning argument for the multiverse is not guilty of it either.

I conclude my contribution by pointing out that the analogies between the different fine-tuning problems can also be used to highlight why the fine-tuning argument for the multiverse, on its own, fails to establish a conclusive case for belief in the multiverse: while we have overwhelming direct observational evidence for the existence of planets beside Earth and for the existence of contemporaries of our ancestors, we do not have such evidence for the existence of other universes; and while we have a solid theoretical understanding of how organisms reproduce and how planets emerge, we have only some speculative ideas, based on inflationary cosmology and string theory, of how universes arise—if they do. While the apparent fine-tuning of our universe would be unsurprising if we had independent osbervational evidence in favour of other universes and a solid theoretical understanding of how universes are generated, the fine-tuning evidence by itself does not fully warrant rational belief in the multiverse.

Roberto Fumagalli. How 'Thin' Rational Choice Theory Explains The proponents of rational choice theory (henceforth, RCT) frequently praise this theory's explanatory potential and allege that RCT provides informative explanations of observed choices (e.g. Becker, 1976, Satz and Ferejohn, 1994). Conversely, many critics complain that RCT applications fail to explain such choices. In particular, several authors (e.g. Alexandrova, 2008, Guala, 2012, Morgan, 2006, Sen, 1987) build on the contrast between so-called 'thick' and 'thin' interpretations of RCT to argue that RCT falls prey to the following dilemma. On the one hand, there is a thick interpretation of RCT, which regards choices as the outcome of a process of instrumental reasoning and rests on empirical assumptions about the neuro-psychological substrates of choice. On the other hand, we find a thin interpretation of RCT, which provides a purely formal axiomatic representation of consistent choice patterns and makes no claim about the neuro-psychological substrates of choice. Thick RCT can be used to explain choices, but is vulnerable to falsifying empirical evidence from neuro-psychology. Conversely, thin RCT is insulated from falsifying empirical evidence from neuro-psychology, but cannot explain choices.

If correct, these criticisms would have far-reaching implications for scientific modellers, since RCT applications figure prominently in a vast range of disciplines (e.g. Boudon, 2003, on sociology, Green and Shapiro, 1994, on political science, Sugden, 1991, on economics). In this paper, I draw on often-cited applications of RCT to demonstrate that contra such criticisms thin RCT can and does explain choices. The paper is organized in two main sections:

- In Section 1, I identify three respects in which I take thin RCT applications to be explanatory and illustrate my thesis with examples from economics and other decision sciences. My main argument proceeds as follows. Thin RCT abstracts away from all information concerning the neuro-psychological substrates of choice. This precludes thin RCT applications from counting as explanatory under various accounts of scientific explanation (e.g. Craver, 2006, Kaplan and Craver, 2011, on mechanistic accounts; Lewis, 1986, Salmon, 1984, on causal accounts), but does not detract from the explanatory potential of thin RCT. On the contrary, the axiomatic derivations at the core of thin RCT provide informative insights as to why agents who differ radically in their neuro-psychological makeup can exhibit choices with the structural (e.g. consistency) patterns defined by thin RCT axioms. These insights, in turn, are explanatory in at least three senses. First, they demarcate the class of actual, possible and counterfactual systems that can exhibit choices with the same structural patterns (structural component). Second, they explicate why many of the properties and features that differentiate real-world agents and the agents posited by thin RCT do not make a difference to these agents' choice patterns (unificationist component). And third, they enable modellers to determine how the choices of real-world agents deviate from the choices of the agents posited by thin RCT under a set of actual, possible and counterfactual conditions (counterfactual component).
- In Section 2, I defend my thesis that thin RCT can explain choices from a series of objections put forward by the critics of RCT. More specifically, I address in turn: the objection from spurious explanations (e.g. Guala, 2012); the objection from causal explanations (e.g. Reiss, 2012); the objection from partial explanations (e.g. Sugden, 2011); the objection from axioms' untenability (e.g. Sen, 1987); and the objection from interdisciplinary consilience (e.g. Craver and Alexandrova, 2008). In addressing these objections, I differentiate my thesis from other authors' accounts of how models that abstract away from empirical information about their targets can be explanatory (e.g. Bokulich, 2009, Bueno and Colyvan, 2011, on the structural component; Batterman and Rice, 2014, Rice, 2015, on the unificationist component; Hindriks, 2013, Odenbaugh, 2005, on the counterfactual component). I then explicate my thesis' implications for the ongoing debate concerning the explanatory potential of RCT and the comparative merits of entrenched philosophical accounts of scientific explanation.

Marion Godman and Martin Bellander. Is there anything I can learn about myself from quantitative psychology? This paper aims to discuss a central problem in the domain of psychology and draw out some methodological and ontological implications that as of yet have not been appreciated in the psychological and philosophy of science literature (cf. Hood 2013). To study a certain process, or a relationship between variables, the majority of studies in quantitative psychology use a large-sample approach. This means measuring a large number of individuals at a specific point of time (or, sometimes, a few different time points). Researchers then analyze the co-variation between the different measurements among these individuals. Say, for example, that we want to find out how episodic memory relates to the amount of social activities. We start by testing a large sample of individual's performance on an episodic memory task and ask them how much time they spend on social activities per week. We calculate the correlation coefficient, and let's say we end up with a value of r = .6. We conclude that social activities are good for episodic memory, and, perhaps, recommend aging people to engage in more social activities.

However, it is entirely possible that every aging individual who follows our advice and increase their engagement in social activities will get a worse episodic memory. This is because the relationship within individuals need not be mirrored by the relationship that exists between individuals. Within an individual we might actually find the opposite relationship. It is possible that social activities are worse for episodic memory than most other everyday activities, so that engaging in more social activities might result in worse memory performance. The positive relationship between individuals might then just be explained by people who have a really good episodic memory also having more friends and therefore engaging in more social activities. The relationship between variables within and between individual would then be opposed. This could have been discovered if individuals would have been measured on multiple occasions, and the correlation coefficient would have been calculated separately for each individual. Moreover, even if the relationship in the present example would be worrisome and easily missed using standard statistical methods, the situation could be even worse. It might be the case that individuals follow an arbitrary number of different models at the individual level.

The possibility of differences in between-person structure and within-person structure are not confined to relationships between directly measurable variables; they could also be true of latent variables used to describe the cognitive architecture of people. To empirically investigate whether this is the case, many participants need to be tested on many occasions over time, which would allow for testing the differences in between and within person structure. This has been done in recent empirical work, demonstrating that the structure of several psychological constructs, such as general intelligence, differs reliably between individuals (Reference excl for review). These latter results also reveal that these differential patterns amongst (latent) variables between and within individuals are not a mere theoretical possibility but in fact empirically verifiable. In light of this, we argue that two different routes for future (quantitative) psychology can be identified. We also aim to begin an evaluation of the merits and drawbacks of each approach.

The first strategy is methodologically radical but ontologically conservative. Its claim is that if the between-individual processes and relations are not reflected within individuals, so much the worse for the existing methodology. Quantitative psychology should then be re-oriented (back) to the individual - the proper home for the discipline (e.g. Hamaker, 2012). The strategy typically recommends that study designs need to be longitudinal, collecting multiple data points for each subject, and that the statistical methods used to analyze the data needs to be adapted to handle this kind of data. For example, time series analyses has been proposed as an important statistical method for future studies in psychology (Hamaker et al. 2005). It is still an open question in this framework whether any of the existing variables typically used in large-sample approaches - let alone latent variables - will survive as part of the ontology of this reformed methodology.

The second strategy is in contrast methodologically conservative and, at least potentially, ontologically radical. It claims that if the successful existing research program of quantitative has discovered important patterns and processes due to its use of large sample sizes it ought to be retained. It may be that many of these between-individual processes are not be mirrored by processes within individuals (though some of course could still be) and so do not tell us anything about individual cognitive processes, but, we should not throw the baby out with the bathwater. Instead of simply eliminating the constructs and latent variables like, general intelligence, from our theoretical apparatus and changing our methodology, we should radically reinterpret these ontologies, for example by understanding the constructs as e.g. multiply realized in individual architectures (Borsboom et al 2003; Borsboom et al. 2009). But while multiple realization may be a correct hypothesis, it only amounts to a negative statement; we still need to understand what explains the robust patterns discovered if not within the individual and this might in turn actually lead us to propose very different ontologies - which we as of yet have little idea about.

Márton Gömöri and Gábor Hofer-Szabó. On the meaning of EPR's Criterion of Reality EPR's famous Reality Criterion (RC) introduced in their 1935 paper on the incompleteness of quantum mechanics is the following claim: If, without in any way disturbing a system, we can predict with certainty (i.e. with probability equal to unity) the value of a physical quantity, then there exists an element of physical reality corresponding to this physical quantity. Purely looking at its wording, it is striking how general and deeply philosophical this criterion is. It is a principle of apparent epistemological character that provides a way, as EPR put it, of ``recognizing a physical reality''. Whenever we are in a certain epistemic position, of predicting the result of a measurement without influencing it, we had better adopt a certain ontology. In other words, the RC can be taken as a general inference pattern from the epistemic to the ontic. In this respect it is on a philosophical par with Quine's and Putnam's ideas about ontological commitment based on a successful scientific theory.

On the other hand, it is remarkable that in the entire history and philosophy of physics---to our knowledge at least---this general epistemological principle has been articulated and applied only once, namely in the EPR argument. There is no mention of anything like the RC outside of the specific context of the EPR problem in quantum mechanics. Moreover, as Fine (1996) points out, while the RC plays a central role in the original EPR text it is completely lacking from Einstein's latter arguments on the incompleteness of quantum mechanics. EPR regarded the RC as a ``reasonable'' principle which is ``in agreement with classical as well as quantum-mechanical ideas of reality.'' It is not entirely clear what Bohr's position was towards the criterion itself, but perhaps it is grounded in his positivistic views that he had doubts about its plausibility, as expressed in his response to the EPR paper (Fine 1996). Refining Bohr's subjectivist conception of quantum theory, Quantum Bayesians have come to claim that the RC is an unwarranted principle (Fuchs, Mermin, and Schack 2014). Most physicists today, however, would probably take it that the RC articulates a kind of ``classical

realism'' that can no longer be maintained in the quantum domain, as is conclusively demonstrated by Bell's theorem (Werner 2014). Others object. Tim Maudlin (2014) recently argued that the RC is far from being a substantial claim; it is rather an analytic truth. As he puts it, the RC is ``just not the sort of thing that can coherently be denied.'' If there is something wrong with the EPR argument or with ``classical realism'' in general, it is certainly not the RC to be blamed for, simply because the criterion is a tautology---so Maudlin. Is then the RC a statement without any content? Might this be the reason for the lack of attention it gained in philosophy of science literature, as well as for its absence in Einstein's latter writings on the EPR scenario?

Or, rather, is the RC is the expression of a substantial philosophical commitment that capable of reevaluation---in particular, capable of reevaluation in light of the results of the EPR--Bell-type experiments? The aim of this paper is to answer these questions and clarify the meaning of the Criterion of Reality. We will argue for the following two claims.

First, the EPR argument and Einstein's latter arguments are different type of arguments with different conclusions. Whereas the EPR argument attacks quantum mechanics from the angle of completeness and shows that all but one interpretations are incomplete; Einstein's latter arguments attack quantum mechanics from the angle of, what we call, soundness and show that the Copenhagen interpretations is unsound (wrong). We will show that what makes the difference between the two type of arguments is the RC itself; hence it cannot be a tautology. The reconstruction of these arguments will be based on a subtle comparison of what elements of reality the various interpretations of quantum mechanics posit on the one hand; and what elements of reality the general metaphysical principles posit on the other. We will see that the RC is just such a principle. Second, it will be argued that the RC is a special case of Reichenbach's Common Cause Principle. The basic idea is a simple translation of the RC to a language describing event types and token events localized in space-time. The main steps of the translation are the following:

``We can predict with certainty (i.e. with probability equal to unity) the value of a physical quantity'' means that there is perfect correlation between the results of our act of prediction and the predicted outcomes of measuring the corresponding quantity;

``without in any way disturbing a system'' means that our act of prediction does not causally affect the predicted outcome event;

``there exists an element of physical reality corresponding to this physical quantity'' means that the outcome of the corresponding measurement is predetermined by some physical event.

Once such a translation is carefully done, the RC essentially boils down to the claim that perfect correlation of causally separated events can be explained by means of a deterministic common cause. Accordingly, it will be demonstrated how the EPR argument can be fully reconstructed from the premises of the Common Cause Principle, locality and no conspiracy. These considerations will lead us to conclude that the Criterion of Reality articulates a deep philosophical principle far from being analytic---at least insofar as the Common Cause Principle can be regarded so.

Leah Henderson. The No Miracles Argument and the Base Rate Fallacy
The no miracles argument (NMA) has been called the 'ultimate argument' for scientific realism (Van Fraassen 1980, Musgrave 1988). Recently however it has been alleged that the NMA is fundamentally flawed because it commits the base rate fallacy (Howson 2000, Lipton 2004, Magnus and Callender 2004, Howson 2013). The allegation is based on the idea that the appeal of the NMA arises from inappropriate neglect of the base rate of approximate truth among the relevant population of theories. In this talk I will argue that the base rate fallacy objection to the no miracles argument fails. Therefore the argument should not be rejected on these grounds.

It has been found in psychological studies that people have a tendency to neglect base rates. A classic case occurs in a medical setting where doctors conduct a test for a very rare disease, which occurs in 1 in 1000 of the population. Suppose the test is such that there is virtually no chance that a person with the disease will test negative (i.e. the false negative rate can be regarded as zero), while there is a small chance of a person without the disease testing positive, say about 5% (i.e. the false positive rate is 5%). A patient takes the test and his result is positive. The probability is actually still quite low (under 2%) that the person has the disease, given the positive test. However when confronted with this problem, people tend to think that the person is quite likely to have the disease, because they neglect the very low base-rate of the disease in the population.

The NMA has been formulated in two different versions, a 'global' version, and a 'local' version. The global version concerns the success of science as a whole. The classic formulation of this argument was first given by Hilary Putnam, and developed by Richard Boyd and Stathis Psillos (Putnam 1975, Boyd 1983, Psillos 1999). The idea is that science as a whole is highly successful, and scientific realism provides the best explanation of that success. Some authors have preferred to formulate the NMA as a 'local' argument (Musgrave 1988, Worrall 2005). In this formulation, the explanandum is the success of a particular theory and the explanans is the approximate truth of that theory. The argument then says that the success of the theory is best explained by the approximate truth of that theory.

The base rate fallacy objection is directed at the local version of the NMA. In the local NMA we assume that it is much more likely for a theory to succeed if it is approximately true, than if it is not. Suppose then that a theory is successful. According to the local NMA, the theory is therefore quite likely to be approximately true. However, it is alleged, this conclusion is reached because of base rate neglect. The complaint is that the proponent of the NMA has failed to take into account of the fact that the base rate of approximately true theories is actually very low. If that is taken into account, the probability that the theory is approximately true, given that it is successful, should actually not be high.

There is however an important disanalogy between cases of genuine base rate neglect and the no miracles argument. In the medical example it is assumed that that the patient in question was randomly drawn from the population as a whole. If this were not true, say because the sampling process is such that those being sent for test are more likely to have the disease, then it is no longer appropriate to take the prior to be the base rate in the overall population. And in this case, it need not be so unlikely that the person does have the disease, if she tests positive.

Similarly, if the scientific method selected a group from the general population which was not just a random sample of all theories that fit the data observed so far, then again it would no longer be appropriate to take the prior to be the base rate of approximately true theories in the overall population of skeptical alternatives. And again, in this case, it is not necessarily unlikely that the theory is approximately true, if it is successful.

The moral here is that one should only take the base rate in the population as one's prior if one knows that the individual was randomly drawn from the population. In the cases which have been presented to subjects in psychological experiments on the base rate fallacy, this is usually stipulated, or at least implied. However, in the case of the NMA, we cannot assume that the theory we have in hand was drawn randomly from the population of skeptical alternatives. In fact, the nature of the sampling process is exactly the point at issue in the realism debate.

On this point, the local NMA does not stand alone, but is supported by the global NMA. According to the global NMA, a high overall proportion of success among theories serves as evidence that the scientific 'sampling' procedure is not random, but rather biased in favour of approximately true theories. That is, the fact that scientists so often succeed in coming up with successful theories is a reason to think that the scientific method generally produces approximately true theories.

The failure of the base rate fallacy objection has more general implications. The base rate fallacy objection has been taken as one reason to turn away from the NMA, in favour of a 'local' or 'retail' approach to the scientific realism debate (Magnus and Callender 2004). However, given that the objection fails, those in favour of a retail approach must turn elsewhere to support their view.

Michael Hicks. When Groups Lie In November of 2015, the New York Attorney General's office opened an investigation to determine whether Exxon Mobil lied about the risks climate change. At issue: did the company believe (or know) that climate change change could hurt its business? Did the company then misrepresent its beliefs to the public? Answering these questions requires us to accept that a company can act and have beliefs, and to acknowledge that in some cases a person speaks for the group rather than for herself. Similar questions arise when we appeal to science to inform technology or public policy--often, individual scientists have views which diverge from the established results of their discipline. When we consult scientists, we should require them to be clear about when they are speaking for the scientific community rather than providing their individual views. When does a scientist speak for herself, and when is she speaking for the scientific community? Answering these questions requires us to provide an account both of group belief, membership in a doxastic community, and group justification.

Most accounts of group belief take the beliefs of a group to be supervenient on the beliefs of group members (Gilbert 1994, Goldman 1999). Alexander Bird (2010, 2014) gives compelling counterexamples to these views, many of which are drawn from standard scientific practice. Bird generates a view according to which whether or not a group knows that P depends on P's place in the group's deliberations, reasoning, and action. Traditional accounts of group belief additionally take group membership to require a shared commitment, explicit or tacit, to adhering to group belief; in response to this, Bird develops a view on which group membership depends instead on participation in the production or consumption of social knowledge.

Here, I show that Bird's view does not give us the tools to represent group belief in addition to group knowledge, nor does it provide an adequate account of group membership. Both are important, as we regularly need to attribute false beliefs to groups to explain or criticize group actions; similarly, we need an adequate account of membership in epistemic groups to delineate responsibility for shared beliefs and to identify suitable group spokespersons. I provide counterexamples to Bird's explicit view and motivate modifying his functionalist account. The view I arrive at is inspired by Bird's, but improves on it by utilizing the notion of a rules constitutive of group functioning. I conclude that a group's beliefs are partially grounded in the rules constitutive of the group, and that group membership requires participation in the rule-governed behavior constitutive of the group.

Finally, I take this view a step further and show how it naturally fits with a process reliabilist framework for group justifiedness. Previous proponents of process reliabilism have focussed on the way in which group beliefs can inherit justification from group members. However, counterexamples to this view are easy to find; nonetheless, an alternative quickly presents itself using the same notion of constitutive rules developed for group belief and membership. I argue that it is the reliability of the rule-governed process of belief formation which matters for group justifiedness; as that process may be independent of the beliefs of group members, the justification of a group belief does not depend on the justification of its member's beliefs.

Laurenz Hudetz. Why supervenience does not imply Nagel reduction
According to the classical view about the relationship between supervenience and reduction, reduction implies supervenience, but supervenience does not imply reduction. More precisely, the classical view has to be regarded as a schema with different instances for the different definitions of the terms `supervenience' and `reduction'. As is well-known there are plenty of them. However, we can broadly distinguish between two main kinds of supervenience and reduction relations, namely those with properties or facts as relata (``ontological level'') and those with theories or vocabularies of theories as relata (``linguistic level''). Relations of the former kind are mainly discussed in metaphysics and philosophy of mind, whereas relations of the latter kind play a greater role in logic and philosophy of science. In this talk, I will focus on an instance of the classical view that concerns supervenience and reduction relations on the linguistic level.

Contrary to the classical view, some philosophers have pointed out that, on the linguistic level, supervenience actually implies Nagel reduction, given certain background assumptions (cf. Kim, 1978; Tennant, 1985; Butterfield & Isham, 1999; Niebergall, 2000; Mainwood, 2006; Butterfield, 2011). However, none of them is particularly happy about supervenience collapsing into reduction (so they usually go on to argue that in most relevant cases the background assumptions do not hold). The main argument in favour of the collapse rests on a prominent result from mathematical logic, namely Beth's definability theorem, which says that implicit definability implies explicit definability. The argument essentially stems back to Hellman & Thompson (1975) and has been brought forward most explicitly and recently by Butterfield (2011). Its basic idea is as follows. Supervenience can be identified with implicit definability and explicit definability is sufficient for reduction. By Beth's theorem, implicit definability implies explicit definability. Thus, supervenience is sufficient for reduction.

In this paper, I will defend the classical view against this argument. I show that supervenience does not imply reduction even under the background assumptions of the main argument sketched above. I pinpoint the reason why there is no collapse. The reason is that explicit definability is not sufficient for reduction. Explicit definability only implies the connectability condition (in the sense of Nagel), but by no means the condition of derivability. However, the latter condition is crucial for reduction, as I shall argue.

This way of vindicating the classical view is novel. So far the main strategy has been to block the collapse of implicit definability into explicit definability established by Beth's theorem (cf. Hellman & Thompson, 1975; Tennant, 1985; Hellman, 2015), usually by arguing that the conditions for applying Beth's theorem are not satisfied in most relevant cases. For example, it has been argued that most scientific theories cannot or should not be formalised in first-order logic (for which Beth's theorem holds), but only in stronger logics that do not have the Beth property (e.g. full second-order logic). In contrast, the vindication of the classical view proposed here does not reject the equivalence of implicit and explicit definability and thus does not rule out that theories are formalised in a logic having the Beth property.

Defending the classical view matters because there are philosophical positions that depend on the classical view. Positions such as non-reductive physicalism are not tenable if supervenience implies reduction. Moreover, one might even regard the classical view as an adequacy condition for the respective notions of supervenience and reduction. Then it should not be violated. It is worth noting that regarding the classical view this way does not trivialise the problem: the main argument for the collapse has been formulated for explications of reduction and supervenience that are by no means trivially inadequate, namely Nagel reduction (or relative interpretability) and implicit definability (Butterfield, 2011). It would be a rather surprising point if it turned out that these explications were inadequate because they violated the classical view. So it is a non-trivial thing to know that the argument contains an error and there is no collapse after all. Thus we do not have to blame the explications.

Defending the classical view in the new way sketched above rather than by denying the applicability of Beth's theorem to actual scientific theories matters for the following reasons. (a) It shows that the additional constraints and assumptions that people have introduced so far in order to defend the classical view are not necessary for that purpose after all. Some of these defence strategies have been criticised as ad hoc (see Kim, 1978). (b) Moreover, the collapse has been taken as an objection against the application of first-order logic in philosophy of science because one background assumption of the main argument is that the theories in question are first-order. My result shows that this is not a good objection against formalisations of scientific theories in first-order logic (such as the project of axiomatising special and general relativity in first-order logic carried out by Andreka, Madarasz & Nemeti (2004, 2007, 2012)).

In this talk I will (1) explain the relevant notions of reduction and definability in sufficient detail and (2) give a precise formulation of the main argument and show what its background assumptions are. Then, (3) I will demonstrate why this argument fails by providing counterexamples from mathematics and physics. Finally, (4) I will deal with possible objections to the proposed vindication of the classical view.

Tero Ijäs. Integrating clinical and nonclinical data in biomedical sciences: a case from cancer genomics The amount of biological and medical data keeps growing exponentially. This expansion of data has since 1990s lead to higher independence of data gathering activities and establishment of new data-intensive biological fields, such as 'omics' (e.g., genomics) fields or later systems biology. Biological research has taken advantage of this abundance of data, and gathering and analysis of large data sets is seen as a fundamental part of understanding biological entities and interactions. However, not all promises of data-driven biology have yet been fully realized and there are more and more attempts to bring Big data perspective to other fields, such as medicine.

This paper explores the question of data integration in biomedical sciences. It especially focuses on the problem of integration of clinical data, such as data concerning biomarkers, body characteristics and health history, and nonclinical data, such as data from animal models, in vitro experiments and omics databases. The paper analyzes the methodological and philosophical challenges that these integration tasks face in their attempts to combine and translate data from multiple disparate and sometimes inconsistent sources. These challenges include the assessment of evidential weight and scope of different (clinical and nonclinical) data sources, special requirements of clinical tools and interventions, locality and sharing of data, tradeoffs of more individualized clinical practice, and interdisciplinary collaboration. The paper focuses on the question of data integration, but explores also the other types of integration, such as methodological and explanatory, as they relate to data (cf. O'Malley and Soyer 2012). Finally, the paper analyzes these problems in relation to cancer genomics.

Sabina Leonelli has analyzed both the nature of biological data (Leonelli 2009) and the role of data integration through the case study of plant biology (Leonelli 2013). This paper follows her analysis on data and typology of integration, but relates these questions to special characteristics of biomedical sciences and data integration in clinical setting. Leonelli (2009) defines data as any material that can be taken as evidence for a range of phenomena, and that can be shared among scientists. As most experimental data is at first local, fruitful sharing and circulation of data requires the scientist to standardize, annotate, and package the data to databases to make it nonlocal. In medicine, where clinical trials are seen as the gold standard of evidence, use of nonlocal database data brings challenges to assessment of evidential weight and scope.

Integration of data-intensive methods to biomedicine is seen as highly valuable and promising for public health. It can help to develop new preventative and therapeutic methods that allow clinicians better to predict and intervene on health histories, or assess therapeutic options for each patient. However, as this paper shows, question of integration is also highly challenging in biomedical sciences. Utilization of large scale genetic data for successful predictions and interventions is difficult in general biology and these problems are amplified in genetic medicine. Translation of genetic data to clinical practices, let alone health policies, might require the combination of basic research and clinical applications, integration of knowledge and expertise from multiple disciplines, as well as practical assessment of safety, utility and costs (Khoury et al. 2012; Wolkenhauer et al. 2013). This has also lead to development of computational platforms to integrate clinical and nonclinical data, as well as emergence of new interdisciplinary biomedical fields, such as translation medicine, personalized medicine (or genomics), and systems medicine. Each of these fields has slightly different goals and methods, but generally they all aim to bring discoveries from (data-intensive) basic science to clinical practice (from 'bench-to-bedside') and consequently, make clinical practice more individualized.

The paper explores these questions through a case of cancer genomics and therapeutics. Cancer provides a good example to analyze the issues of data integration in biomedical sciences, as it is both valuable and challenging problem. Cancer is a highly complex and heterogeneous disease (or rather, a group of diseases), and no general theory can capture the causal dynamics of every cancer case (Plutynski 2013). Individual cancer cases are highly variable and conflated with individual's genetic and environmental factors, which complicates the choice of appropriate therapy. Importance of correct prediction and choice of therapy is amplified, as most cancer therapies are highly invasive and stressful to patients, and costs of ineffective therapy are high. Molecular knowledge of cancer dynamics is not yet transformed to successful clinical applications, and therefore personalized medicine perspective is seen as a promising avenue of research in cancer genetics. This individualization of cancer therapy elaborates the general problems of data integration, locality of data and tradeoffs between generality and realism analyzed in this paper.

Elizabeth Irvine. Evolutionary Reasoning in Studies of Language Origins Studies of the origins of symbolic language are hampered by a lack of direct evidence, and so regularly use an explicitly interdisciplinary and multi-method approach to converge on robust models and explanations of the target phenomenon. However, there is rarely direct attention paid to identifying relevant selection pressures, and the adaptive landscape they work within, that would go towards explaining the emergence of linguistic symbol systems.

This is potentially highly problematic in terms of evolutionary reasoning, as symbol systems will not emerge when there is no pressure for them to do so, when the costs of setting up such systems are not outweighed by their benefits, and when those costs (for whatever reason) cannot be paid. For example, while it is often suggested that it appears useful, and perhaps even necessary, for hominids to use symbols to communicate about hunting parties or tool making or other physical skills, it may well be the case that other non-symbolic means of communication are sufficient. Accordingly, this paper first presents some theoretical arguments against the need for symbols for a range of communicative needs, and second some experimental evidence that throws light on the adaptive landscape relevant to setting up new forms of communication.

First, non-human primates can point (Lyn et al. 2011; Moore et al. 2015) which strongly suggests that pointing was among the repertoire of hominid communication, and theoretical work suggests that pointing and non-referential gestures can be incredibly powerful means of communication. Pointing can be used fairly simply to direct a conspecific's attention to an item in the immediate environment, or it can be used to refer to absent objects or distant places, or to convey complex messages about shared experiences, values, or plans. A point is merely a way of directing attention to a (often vaguely specified) location in space, so communication using pointing is largely done via mind-reading capacities (e.g. inferring the communicative intentions of a conspecific) (Tomasello et al. 2007), which in turn makes use of shared common ground.

Depending on levels of shared common ground and mind-reading abilities, pointing can then potentially do a huge amount of communicative work without the need for symbols. The meaning of a point will be particularly easy to infer where communication occurs within heavily routinized activities, which likely extends to complex tool use and tool preparation, and group hunting, where shared skills, goals and experience can be used to interpret pointing gestures (and other non-symbolic forms of communication).

Second, recent experimental research (Irvine & Roberts 2015) suggests that (adult human) participants completing reasonably complex co-operative tasks often use a range of ways to communicate, but do not develop symbol systems, even where it is more efficient to do so. Instead, participants use pointing, develop efficient task strategies to minimize the need for communication, and also develop a non-referential and highly fluid means of communicating about sequence organization (e.g. gestures for 'that's right, move on', 'no, wrong', where sometimes the same gesture can be used for both). The costs for setting up a symbol system either cannot be paid by novices at the task (e.g. due to lack of time or co-ordination), or are not worth paying for those who are proficient at the task, so in both cases they go unpaid. If contemporary humans fail to set up symbol systems while doing these kinds of tasks, this suggests that less symbol-savvy hominids were even less likely to.

Of particular importance here is that the costs and problems in setting up a symbol system are rarely modeled but are very real: 'egocentric bias', or the failure to take the other person's point of view regularly leads to failure in communication games (Galantucci et al. 2012), and taking time out from a task to set up a convention is directly costly. What these experiments further show then is that not only are conventions hard to set up for novices at a task, but once participants are skilled at a task, there is even less reason to set up a symbol system, as other means of communication are sufficient (indeed are often highly efficient). Further experimental work in this area will help to outline just when symbol systems are worth paying a start-up cost for, when they can be paid, and how inventive participants can be in using pointing, other non-representational gestures, and efficient task and communicative strategies, before they set up a symbol system.

Indeed, one possibility here is that symbolic communication is only very rarely worth paying a cost for in order to plan and monitor co-operative physical activities among (skilled) individuals. However, symbolic communication may be fundamentally necessary to enable the regulation of long-distance and long-term reciprocal cooperation, allowing language origins to be placed in line with other social innovations in hominid lifeways (e.g. Sterelny 2015). Overall though, further theoretical and experimental work is needed to identify the selection pressures and adaptive landscape relevant to symbol use, in order to inform evolutionary reasoning about the origins of symbolic language.

Milena Ivanova. Poincaré on Beauty in Science While Poincaré's views on the aim of science have received significant attention, his views on the role aesthetic considerations play in scientific practise have remained unexplored. In this paper I offer a systematic new analysis of Poincaré's understanding of beauty in scientific theories and show how this account fits with his views on the aim of science.

Aesthetic judgements are integral part of scientific practise. Scientists employ aesthetic judgements in the selection of phenomena to study, the construction of hypotheses, the evaluation of theories and in deciding their epistemic commitments towards a theory. For Poincaré beauty is an important motivation for the study of nature. He argues that "[t]he scientist does not study nature because it is useful to do so. He studies it because he takes pleasure in it, and he takes pleasure in it because it is beautiful" (Poincaré 2001, 368). Poincaré clarifies that beauty equates to the harmonious order that our theories reveal and not to the beauty that 'strikes the senses'. Poincaré concerns himself with the beauty found in our theories that unveil unity and harmony in the phenomena. According to him the aim of science is to offer us understanding of the underlying relations between phenomena, of the harmony in nature. It is in this underlying harmony or unity that our theories uncover that we find beauty. I argue that for Poincaré beauty is an aesthetic property that reduces to the elegance and unity of scientific theories. I reconstruct the argument for simplicity and the argument for unity in order to show their relationship and the overall consistency of Poincaré's position.

Starting with the argument for simplicity, I address the following three questions: (1) how is simplicity defined; (2) how is it used; and (3) how is it justified. I argue that Poincaré is mainly concerned with the mathematical elegance of scientific theories. When it comes to the applicability of simplicity, I argue that Poincaré takes simplicity to play a purely heuristic, not epistemic, role. That is, simplicity does not lead us to true theories; rather, it aids our decision-making. Aesthetic values guide our choice in the construction and selection of hypotheses (ibid., 99), but these are not regarded as objective properties. Whilst simplicity plays a heuristic role it does not have epistemic significance, it leads to convenient not true theories. I support this claim with what I call the 'historical argument', offered by Poincaré, which shows that the development of science sheds doubt that nature in itself is simple (ibid., 99-100). He argues that simplicity is not linked to truthlikeness, it is not a guide to the true nature of reality, but is a condition of our making.

Turning to unity, Poincaré identifies unity with the grasp of the harmony between the phenomena that scientific theories give us, which ultimately provides us with understanding of these phenomena. I argue that the harmony our theories reveal cannot be understood in either objectivist of projectivist terms. It is not an objective feature of the world outside our mental capacities; however, it is not a merely subjective feature projected on nature by us. Appealing to a form of intersubjective validity, Poincaré argues that unity is part and parcel of our intellectual capacities and an ideal we follow in our enquiries (ibid., 396-397). Unity is a guiding principle in the selection and evaluation of scientific hypotheses (ibid., 368). For Poincaré, beauty is to be found in the harmony our theories reveal. It is to be found in the hidden relations that our theories uncover and the unification that they give us in showing how different, apparently disconnected phenomena, relate. It is this harmony that Poincaré takes to give us understanding.

I argue that for Poincaré simplicity and unity are regulative ideals that needed to be followed because they are linked to the ultimate aim of science – gaining understanding of the relations that hold among the phenomena (ibid., 112). For Poincaré, beauty is to be found in the harmony our theories reveal. It is to be found in the hidden relations that our theories uncover and the unification that they give us in showing how different, apparently disconnected phenomena, unite. It is this harmony that Poincaré takes to give us understanding.

To further defend my account of Poincaré's aesthetics of science, I analyse the link between aesthetic judgement and utility. I turn to Poincaré's account of creativity in scientific discovery. It is here that Poincaré makes the important link between utility and aesthetic judgement. He argues that it is our aesthetic sensibility that guides the selection of useful and fertile hypotheses during the creative process. Poincaré argues that the useful ideas are the ones that trigger the scientists' aesthetic sensibility. It is in this context that he appeals to simplicity and harmony. The aesthetic sensibility, Poincaré argues "plays the part of the delicate sieve" which checks the result blindly generated by the mind and selects only the most elegant and beautiful combinations produced (Poincaré ibid., 397). For Poincaré the aesthetic sensibility selects the theories that best suit our aesthetic requirements, but he also claims that "[t]he useful combinations are precisely the most beautiful" (ibid.).

I argue that by reducing aesthetic judgements ultimately to being judgements about the unity and simplicity of scientific theories, Poincaré offers an interesting reductivist account of aesthetic properties. I show the link between utility and aesthetic sensibility and argue that for Poincaré beauty is indicative of understanding rather than truthlikeness. I show that Poincaré's account does not easily fall between the projectivist or objectivist views in aesthetics and also departs from a strictly realist position. I show how this view of beauty of scientific theories and creativity are compatible with his views that the aim of science is the construction of a convenient system of relations that aims to describe the structure of nature.

Anneli Jefferson. Mental disorders and physical disorders - an obsolete distinction? The search for physical causes of mental disorders has a longstanding tradition in psychiatry, and attempts to find the so-called biological basis of mental disorder are currently wide-spread in the biomedical sciences. Some form of mind brain identity is generally accepted by most scientists working on the mind and the brain, but there is less agreement what implication, if any, this has for the concept of mental disorder. Particularly the question whether physicalism implies that mental disorders are best understood as resulting from, or expressive of, problems at the level of the brain is contested.

In my talk, I consider three reasons why we might be sceptical of the claim that the category of mental disorder and brain disorder might be merged into one. The argument from multiple realizability is frequently cited as a reason why attempts to reduce mental disorders to brain disorders are likely to remain unsuccessful. One type of mental state or function, so the argument goes, can in principle be realised in different ways by the brain. Thus, we cannot count on being able to find the same brain differences across individuals with one and the same psychological diagnosis. In other words, we may not find specific brain differences that underlie a given mental disorder. The extent to which mental problems de facto are associated with specific differences in brain structure or function is an empirical question. At present, the jury is still out on whether we are likely to find brain dysfunctions that correspond with mental dysfunctions in a significant number of cases. This issue is further complicated by the fact that it is not clear whether current diagnostic criteria for mental disorders cut psychological reality at its joints, i.e. whether symptom based nosologies have the correct level of grain or are too broad or too narrow. So, while multiple realizability creates conceptual space for mental disorders which are not at the same time brain disorders, the extent to which these categories will converge is an empirical question.

The second point I consider draws on an analogy between the mind/brain and the hardware/software distinction. Philosophers have pointed out that just as the existence of a software problem does not necessary entail the existence of a hardware problem, the existence of mental problem need not entail the existence of a brain problem (Papineau 1994, Boorse 1976). There could be mental dysfunctions which are not at the same time brain dysfunctions. I concede this point but point to the fact that our understanding of what constitutes a brain problem or malfunction is at least partially dependent on the association between specific brain differences and specific mental dysfunctions. While there are some clearly identifiable cases of brain dysfunction or damage, such as lesions matters are not always so clear cut. For instance, dysfunctions such as amygdala hypo-function are diagnosed with reference to the psychological dysfunctions associated with these differences in brain function. The individuation of brain dysfunction therefore becomes partially dependent on that of mental dysfunction. This allows for the theoretical possibility that our understanding of brain dysfunction becomes increasingly psychologised.

Finally, I discuss a thought experiment by Hanna Pickard (2009), which aims to show that the presence of mental dysfunction is essential to our understanding of a condition as a mental disorder, and that our conception of mental disorders is therefore distinct to that of brain disorders. Pickard asks us to imagine that we have found the neurological basis of schizophrenia which was reliably present in cases of schizophrenia. She then presents us with the case where there is an individual with the neurological profile associated with schizophrenia, but without the psychological problems which we currently use to individuate schizophrenia. She argues that under these circumstances, we might well be happy to say that the individual has schizophrenia, but would deny that they are mentally ill. This argument establishes that we currently see some kind of mental dysfunction or problem as a necessary condition for calling somebody mentally ill.

I compare Pickard's thought experiment with a pertinent real-life comparison case and draw out some relevant implications for our concepts of mental disorders and brain disorders. It is well-established that there are cases of asymptomatic Alzheimer's disease, where people present with the neurological anomalies of Alzheimer's but do not suffer cognitive impairment (cf Driscoll and Troncoso 2011). In cases of asymptomatic Alzheimer's, too, we would not classify individuals as suffering from dementia or mentally ill, even though we would characterise them as having Alzheimer's. Nevertheless, we may still think of them as ill. I suggest that in as far as we have independent reasons for thinking of the neuritic plaques and neurofibrillary tangles as pathological, we will and should think of these individuals as suffering from a biological illness without the attendant mental illness, dementia. However, in cases where our conception of what constitutes a brain dysfunction is solely dependent on the psychological dysfunction with which it is associated, we should not think of that individual as mentally ill if they exhibit the brain anomalies without the psychological anomalies.

I conclude that while our conception of mental disorders and disorders of the brain to become are likely to become increasingly meshed, there are clear conditions under which relabeling mental disorders as brain disorders would be theoretically fruitless and unmotivated.

Ciprian Jeler. Frameworks for understanding natural selection and their impact on multi-level selection theory This paper identifies a number of different characterizations of natural selection and their consequences for multi-level selection theory. It begins by distinguishing between views that keep natural selection and evolution by natural selection apart and views that tend to conflate the two notions. A prominent example of the first view is provided by John Maynard Smith. When he distinguishes between units of selection (i.e. entities that exhibit phenotypic differences that cause fitness differences) and units of evolution (i.e. units of selection whose phenotypic trait of interest also exhibits heritability), he is sharply distinguishing selection from evolution by selection. This is perfectly in line with the distinction between phenotypic selection and response to selection that is commonplace for quantitative geneticists and animal breeders. Importantly, this view only involves one type of entity, one that needs to exhibit variation in phenotype and fitness for selection to occur, and whose offspring must tend to resemble it if there is to be response to selection.

However, a population-genetics tradition inspires a view of selection that does not draw a sharp line between selection and evolution by selection. A good example is the replicator/interactor framework, in its two versions defended by Richard Dawkins and David Hull. While sometimes the same biological entity can fulfill both the roles of replicator and interactor, the roles themselves remain distinct: replicators pass on their structure largely intact in replication, while interactors directly interact, as cohesive wholes, with their environment. However, Hull's definition (from 1980) of the interactor goes on to add that they interact with the environment "in such a way that replication is differential", which seems to indicate that their interaction with the environment is only selectively relevant if it leads to differential multiplication of replicators. This is precisely Dawkins'(initial) position: his interactors (termed "vehicles") are just temporary aggregations of interacting replicators, and therefore the "success" of these vehicles needs to be defined, and not just measured, in terms of differential replication of the replicators that gave rise to them. And since, by definition, replicators are transmitted with high fidelity from one generation to the next, "heritability" thus becomes absorbed into the very notion of selection. The distinction between selection and response to selection is completely erased. (The fact that he equates, in 1978, "replicator selection" with "replicator survival" is a striking illustration of this.)

Hull adopts a more nuanced position when he goes on to define, in the same 1980 paper, selection as "a process in which the differential extinction and proliferation of interactors cause the differential perpetuation of the replicators that produced them". The first part of this definition – the differential extinction and proliferation of interactors – corresponds to what quantitative geneticists call phenotypic selection. However, the latter part of the definition states that selection can only bear this name if the differential replication of interactors also leads to differential perpetuation of replicators. Thus, the success of an interactor seems to be defined as being, at the same time, differential fitness at the level of the interactor and at that of the replicator.

This paper argues that, depending on which of these three views of selection we embrace, there are three possible stances one can take with respect to multilevel selection cases known as MLS1. According to the terminology proposed by John Damuth and Lorraine Heisler, in MLS1 scenarios we are interested in explaining the change in frequency of a particular individual trait in a metapopulation divided into groups. Therefore, in MLS1 scenarios, a group is fitter than another one if it produces more offspring individuals (whereas in MLS2 scenarios a group is fitter if it produces more offspring groups).

If we assume a view of selection akin to that of Maynard Smith or of quantitative geneticists, we have to conclude that MLS1 should not count as group selection. According to this view, in order for selection to take place, phenotypic variation and differential fitness must occur at the same level: an entity cannot be said to be fitter than another one just because it produces more entities of a different, lower-level type. (Inexplicably, in his 2006 book, Samir Okasha explicitly embraces such a view of selection, yet keeps discussing MLS1 scenarios as if they involved genuine group selection).

Ironically, if we embrace a view of selection akin to Dawkins' replicator/vehicle framework, the idea that MLS1 cases do involve group selection becomes more tractable: a group can be said to be fitter than another one if it produces more lower-level units, just like an aggregation of interacting replicators can be said to be more successful than other aggregations if its output of replicators is larger. Of course, having to embrace a Dawkins-like view in order to defend group selectionist interpretations of MLS1 scenarios would probably be enough to deter most group selectionists. If that is not enough, Dawkins' framework has received serious criticism over the years and it is far from certain that embracing such a view remains a sensible option.

On the other hand, Hull's version of the replicator/interactor framework resists some of the stronger objections to Dawkins' framework and, modulo some amendments, might still be embraced (both points have been argued for by Peter Godfrey-Smith, for example). But embracing such a view would force us to reconsider the way we tend to conceive of MLS1 scenarios nowadays: these scenarios would have to be redefined as scenarios in which groups differing in character produce different numbers of offspring groups and, at the same time, also produce different numbers of offspring individuals (i.e. different numbers of the lower-level entities composing the groups). This is consonant with the way in which Michael Wade's famous 1970s experiments were conducted (using a "propagule" mode of reproduction of groups in which a group's production of propagules is proportional to its output of individuals). This is also the way in which Robert Brandon conceived of group selection in his book Adaptation and Environment. The final part of this paper discusses Brandon's arguments and outlines the difficulties of such a view.

Donal Khosrowi. Getting Serious about Shared Features In Simulation and Similarity, Michael Weisberg (2013) offers a similarity-based account of the model-world relation, i.e. the relation in virtue of which scientific models exhibit their predictive and explanatory capabilities. This account accommodates several prevalent intuitions: that similarity comes in degrees; that only relevant instances of similarity matter; and that contextual factors such as modeling goals play an important role in the model-world relation. Moreover, by offering an explicit analysis of similarity, Weisberg addresses the pertinent criticism that similarity-based accounts are uninformative unless they clarify what it means for models to be similar to their targets (see e.g. Teller 2001:399). To address this call for clarification, the central idea that Weisberg offers is that models are similar to their targets in virtue of sharing features with them.

I argue that Weisberg fails to give a successful analysis of similarity as the relation in virtue of which successful models are successful, because he does not offer an adequate account of what it means to share features. More specifically, I argue that shared features can be construed in at least three substantively different ways, each of which creates undesirable consequences for Weisberg's analytical aim.

The first construal I consider is to say that shared features are identical features. This construal is implausible because models are rarely identical to their targets in any interesting ways (cf. Parker 2015). Moreover, I argue that this construal renders Weisberg's account vulnerable to the same criticisms that he offers against rival, isomorphism-based accounts.

The second construal says that features are shared to the extent that they are sufficiently close on some subvenient scale, e.g. when a parameter or prediction is within some threshold range of an estimated or observed value from the target. Even though Weisberg seems inclined to adopt this construal (cf. Weisberg 2015), I argue that it is unlikely to be helpful in realizing his analytical aim. First, the construal can only handle quantitative features, i.e. those that are in principle amenable to quantitative operationalization. This severely constrains the kinds of features that the account can accommodate. Second, for this construal to be informative it needs to specify how thresholds of sufficient closeness are determined. Modeling goals and background theory appear to be plausible candidates for telling us how close, say, a model's parameter needs to be to an estimated value from the target in order for the model to be epistemically successful. However, sophisticated background theory and sufficiently precise modeling goals are often unavailable. In fact, the lack of background theory is often one of the reasons for why we engage in some modeling activity in the first place. Without background theory that is sufficiently developed to determine thresholds for shared features, we are unable to tell whether a given feature is shared. This is a problem at least insofar as Weisberg maintains that his account of similarity must be computable.

The third construal of shared features I consider says that features are shared to the extent that they are sufficiently similar to each other. This construal severs Weisberg's analytical aim because it tells us that models are similar to their targets in virtue of being similar in many important respects and not too dissimilar in too many other important respects. While this sheds some light on how the overall similarity of a model to its target hinges on specific feature-level similarities, it fails to elucidate what constitutes these similarities as well as why they are conducive to models' epistemic success. Even so, while the sufficient similarity construal seems to sever Weisberg's analytical aim, I also argue that at least in some cases and for some kinds of features, shared features are best understood as sufficiently similar features. I offer a test-case from the financial economics literature to substantiate these claims.

Based on these results I conclude that neither of the candidate construals is likely to be successful in helping Weisberg deliver an informative analysis of similarity that does not bottom out with an unanalyzed notion of shared features.

Against the background of this negative result I offer a proposal to revise Weisberg's account as a more general feature-sharing account. More specifically, this revision turns Weisberg's account upside down by saying that the relation in virtue of which successful models are successful, at the most general level, is one of feature-sharing, rather than similarity. On the level of specific features, this revised account says that shared features can be realized by different types of feature-level relations between models and targets, i.e. identities, similarities, sufficient quantitative closeness, isomorphisms, and possibly others. So it is flexible with respect to what shared features consist in, as they may be constituted by any of these relations. Understood in this way, more fundamental questions about whether and why these relations are conducive to models' epistemic utility may be delegated to subvenient accounts of these specific relations. Moreover, in line with Teller (2001) I suggest that this yields a highly local account of the model-world relation in the sense that it depends crucially on information delivered by concrete modeling contexts. Yet, pace Teller, this does not preclude the account from being general in scope. Finally, I argue that this picture allows us to consider purportedly rival isomorphism- and similarity-based accounts as complementary accounts of different feature-level relations that may simultaneously obtain between a given model and its target.

Eleanor Knox and Alex Franklin. Emergence without limits: a case study Recent literature on inter-theoretic relations in physics (e.g. the work of Batterman and Butterfield), has paid a great deal of attention to the importance of asymptotic limiting relations between physical theories. There is a growing consensus, even across the divide between discussants, that such relations signal a kind of emergence, where emergence is meant less strongly than in some literature outside the philosophy of physics. For a phenomenon to be emergent in the physicist's sense is for it to be "novel and robust". Spelling these out, in particular explaining how novelty might exist even where reduction looks plausible, is a challenging and largely incomplete project. One way of reading the literature on asymptotic limits is as picking out one relevant kind of novelty.

But merely looking at the above literature might leave one with the impression that asymptotic novelty is the only, or at least the most important, game in town. In this paper, we'll examine a phenomenon that seems to have many of the hallmarks of emergence, but that doesn't require asymptotic reasoning for its derivation. Our case-study will be phonons – quasi-particles in a crystal lattice. These are occasionally mentioned in the literature (for example in Wallace, [2001]) as an example of emergence. They are a robust phenomenon and explanatorily useful. But perhaps most importantly, the relationship between phonons and the crystal lattice in which they live has much in common with the relationship between the particles of the standard model and the underlying quantum field theory (see Wallace [2001] for more details of the analogy). It therefore seems likely that we can learn something about the emergence of particles by examining the phonon case.

We will argue that phonons are interesting not because they require asymptotic limits for their derivation (they don't), nor even because they describe the system on a different scale from that of the atomic lattice. The sense in which moving to the phonon description involves rescaling is not a particularly interesting one; true, phonons are a delocalised phenomenon and atomic vibrations are a local one, but rescaling is not the essential operation here. Rather, we move to the phonon description via a particularly interesting change of variables involving first a move to a reciprocal space description, and then a series of approximations. The result is a description that, for certain purposes, possesses more explanatory power than the underlying description. It articulates relevance relations by exploiting descriptive redundancies at the atomic level.

Our account of explanatory power here follows that of Knox [2015]. The right change of variable can sometimes facilitate an explanatory abstraction in such a way that the higher-level description may be said to have novel explanatory power despite the relationship with the lower level description being well understood. Phonons illustrate this example nicely, and most of the talk will be dedicated to fleshing out the phonon story. However, it's also important that not just any such variable change begets increased explanatory power; we'll also look at a toy example of coupled oscillating masses on springs (normal modes) that falls short of the kind of descriptive change involved in the phonon case.

With these examples in hand, we will argue that the focus on limiting relations has been distracting and that one can identify putative cases of emergence while sidestepping such issues.

Jonny Lee. Inference to Representation: Scientific Explanation & Two Kinds of Eliminativism Mental representation features heavily in scientific explanations of cognition. The principle of 'inference to the best explanation' (IBE) says that we ought to believe that our best theories are true. If our best theories feature representations then we ought to believe those ascriptions are true. Nevertheless, eliminativism about mental representation remains popular. In its various forms, eliminativism holds that we ought to reject talk of representation in some or all of cognitive science. How do we reconcile the fundamental role of representation in cognitive science with the appeal of eliminativism? In beginning to answer this question we must understand the different motivations belonging to distinctive kinds of eliminativism. This paper will argue that there is an important distinction between a priori and a posteori eliminativism about mental representation. I will begin by outlining the principle of IBE before distinguishing between the two kinds of eliminativism. Next, I will survey some objections to a priori eliminativism, concluding that it is far from clear how we are to interpret this position in the face of central assumptions at the heart of cognitive science. I will finish by considering the possibility that the ontology of representation could turn out to be more nuanced than philosophers have previously allowed.

Our ordinary understanding of representation is that of a state, structure or activity possessing 'content' with truth conditions. As such, purported representations must possess the capacity to succeed or fail to 'be about' something. However, such 'semantic properties' are paradigmatically fixed by social norms. The truth conditions of a sub-personal representations cannot be fixed by social norms though, because they are not the subject of conscious deliberation. Despite the prevalence of explanations in terms of representation, there is no consensus on how to understand what it is to be a mental representation, no agreed upon account of how to naturalise representational content.

How then should we think of mental representations? Despite the aforementioned puzzle, one obvious and persistently popular answer is that representations are real, in so far as ascriptions of representations within cognitive science are (at least mostly) true. Many theorists think that the entities that feature in our best scientific explanations are a good indication of what is true, or put differently, that they are an effective indication of what we ought to believe. This intuition is at the heart of IBE. This principle is fundamental to scientific inference and is the most powerful tool in the argument for realism about mental representation, making a crucial link between representations as a feature in our scientific explanations, and the justification for our belief in representational ascriptions.

Despite the apparent explanatory success of mental representations, and the prevailing appeal of realism given the entrenched appeal of IBE, two distinct forms of eliminativism remain popular. A posteriori eliminativism claims that representations do not in fact feature in (at least some) of our best explanations of cognition, and because of IBE, we should reject the truth of representational ascriptions. A priori eliminativism on the other hand, claims that representations cannot feature in our best explanations of cognition because the properties of representation are logically exclusive from the properties of sub-personal cognition. The former eliminativism accepts representational explanation in principle, but rejects its use in practice, although typically such theorists are 'local eliminativists', seeking to replace representation talk within a particular domain. The latter eliminativism holds that to ascribe representations to sub-personal phenomena is to make a category error. Radical enactivists provide an argument for a priori eliminativism when they say that explanations involving semantic properties are, as a matter of principle, only appropriate for a restricted domain of sophisticated cognition, roughly, cognition involving language or social norms (Hutto & Myin, 2013).

A priori eliminativism faces two related objections. Firstly, it must explain why, given IBE and the plethora of representational ascriptions, we should reject those ascriptions as true. Secondly, it must explain why we should not hold out hope for a future theory of representation, one which can successfully provide a naturalistic account, and accommodate the semantic properties that concern them.

Fictionalism about representations is also an option. For the fictionalist, representations should remain a part of our scientific vocabulary, but not because ascriptions of representations are true. For the fictionalist, statements such as 'neural pattern y represents x' are false, but serve an epistemic purpose (Sprevak, 2013). Fictionalism, however, faces a potent problem: how to reconcile its rejection of the truth of representation with IBE.

It is important to note that fictionalism and eliminativism provide two ways to talk about representations, whilst rejecting their truth. In the former case, representation talk provides some useful role in explanation. In the latter case, representation talk (may) provide a 'temporary gloss', or a short hand pending a further true explanation. But such a gloss is a dismissive acknowledgement of our (at least current) epistemic limitations, not an endorsement (Dennett, 1998). Acknowledging the existence of different ways of treating representational talk, philosophers should allow for the possibility that ascriptions of representation in different areas of cognitive science may be subject to different ontological assessments. The modest combination of IBE, and the continued success of representation within cognitive science, does lend tentative support to the truth of some representational explanations. Nevertheless, we should remain open to the possibility that some representation talk is a mere gloss or entirely inappropriate, the result of overextending an otherwise good concept, and that still other representation talk is a useful fiction, untrue but required to satisfy human epistemic needs.

Moving forward, we must be careful not only to distinguish between the two sorts of eliminativism, in order to properly understand the conceptual space, but to be aware of the various nuanced possibilities for the ontological landscape of mental representations, one that may lie between inflexible and absolute theoretical abstractions.

Dennis Lehmkuhl. The problem of motion in general relativity The general theory of relativity has two equations at its core: the Einstein field equations, which describe the dynamics of gravitational fields, and the geodesic equation, which is the equation of motion of test bodies subject to gravitational fields. The problem of motion, the query of whether the equations of motion can be derived from the gravitational field equations, has been one of the most important questions both for the foundations of general relativity and for its application to astrophysics. Up to now, philosophers of physics have been concerned merely with one of the two major research programmes aimed at accomplishing such a derivation. They have dismissed the second programme, pioneered by Einstein and Grommer in 1927, as being misguided. However, I will demonstrate that the historical development of this programme shows us that it is closely linked to the search for exact solutions to the gravitational field equations. This, in turn, allows us to link the problem of motion to the dynamics of black holes advanced since the 1960s. We will see that the careful interpretation and conceptual analysis of equations of motion and exact solutions to the gravitational field equations allow for an entirely new perspective on the foundations of general relativity.

Bihui Li. Solutions in Constructive Field Theory To date, philosophers of quantum field theory (QFT) have paid much attention to roughly two kinds of QFT: Lagrangian-independent approaches to algebraic QFT, and perturbative Lagrangian-based QFT. Comparatively less attention, however, has been paid to constructive QFT, an approach that aims to rigorously construct solutions of QFT for particular Lagrangians and Hamiltons, ensuring that such solutions satisfy certain criteria. Since we usually take solutions in physical theories to describe the behavior of systems falling under the theory, constructive QFT deserves philosophical attention. I argue that once we look at constructive QFT, we can see that it crucially relies on information from perturbative QFT, in a way that suggests axiomatic QFT is not sufficient to define possible systems in QFT.
Constructive QFT is rather different in approach and aims from axiomatic QFT. While the term 'axiomatic QFT' is used to describe approaches to QFT that aim for theorems about general structure without reference to particular Lagrangians or Hamiltonians, constructive QFT, in contrast, refers specifically to attempts to rigorously construct models of QFT corresponding to particular Lagrangians or Hamiltonians used in mainstream theoretical particle physics. Constructive QFT may be carried out in the algebraic tradition, in the functional integral tradition, or in both. The terms "algebraic" and "functional integral" refer to the mathematics used, with the former using C* algebras and the latter using functional integrals. I focus on constructive QFT in the functional integral tradition, with specific attention to its interaction with perturbative Lagrangian-based QFT.
I examine two aspects of constructive QFT in the functional integral tradition that are relevant to our understanding of the theoretical structure of QFT. The first is the question of what counts as a solution to a specific Lagrangian in constructive QFT. I argue that the criteria for what counts as a solution include some kind of correspondence with perturbation series derived in perturbative QFT. Presently, the form of the correspondence is postulated to be some kind of asymptoticity (in the sense of divergent asymptotic series) rather than convergence. However, the nature of the asymptoticity is disputed, with some proposing stronger or weaker versions of asymptoticity. This correspondence is a kind of "physical criterion" that is not specified in the Wightman axioms or Osterwalder-Schrader axioms, which are the usual axiom systems the functional integral tradition aims to satisfy. Furthermore, the specific mathematical form of this criterion—whether asymptoticity is the correct criterion, and if so, what kind of asymptoticity---is a matter of debate and not something that is simply given by the axioms of QFT.
The second aspect of constructive QFT in the functional integral tradition that I examine concerns the information that constructive field theorists use to construct solutions. I show that they make heavy use of perturbative QFT, and furthermore, that the specifics of the regularization methods, counterterms, multiscale expansions, and so on that are used affect the success of the construction. Constructive field theorists admit that impossiblity theorems in constructive field theory are hard to come by, because in order to prove that a construction of a certain Lagrangian is impossible, one must show that all possible methods of constructing it, namely using all allowable kinds of regularizations, counterterms and so on, would fail. On the other hand, to prove that a construction is possible, one only needs one method of construction to work. In short, a successful construction relies on information that is not just in the axioms of QFT. Since solutions to specific Lagrangians of QFT describe the possible behavior of systems with those Lagrangians, solutions ought to be part of the data for any interpretation of QFT. Therefore, the suggestion is that axiomatic QFT under-constrains the nature of QFT, and perturbative and constructive QFT are worth some philosophical attention.
Finally, I address the possible objection that the two aspects of constructive QFT discussed above concern merely the epistemic process of finding a solution to models of QFT, and do not reflect anything ontological about solutions in QFT. First, I argue that the condition of correspondence to perturbative solutions is not a mere epistemic criterion but a constitutive one. Second, I argue that the perceptual analogy that might tempt one to label the considerations I describe as "epistemic" does not work. In other areas of mathematics, such as differential equations, one may label a method "epistemic" if it is used out of convenience to obtain a solution that is already described by a closed-form formula. However, in many cases we do not have such solutions. Our only way of describing the solutions may be by some approximate expression, such as a series expansion. In such cases, it becomes less clear that the series expansions is only a mode of epistemic access to the solution, rather than a direct expression of the solution itself. I argue that the case of solutions in constructive QFT is directly analogous to the case of differential equations that can be solved only by some approximate method. The perceptual analogy fails in these cases because there is no reason to believe that there is some Platonic solution out there other than the solution that constructive QFT provides, with counterterms, limits, regularizations, and all the gory details provided.
In short, by paying attention to the constructive QFT which is another kind of rigorous QFT distinct from purely axiomatic QFT, we see that perturbative QFT plays a more foundational role in QFT than commonly believed. Perturbative QFT is not merely a dispensable approximation tool, but is essential to constructing and defining solutions in QFT even when we restrict ourselves to rigorous mathematics.

C.D. McCoy. Epistemic Justification and Luck in Inflationary Cosmology I present a recent case in theoretical cosmology, and argue on its basis that explanatory considerations play a crucial role in epistemically justifying theory choice. Much of the philosophical debate over whether explanatory power is an epistemic theoretical virtue has centered on general philosophical considerations, for example underdetermination arguments and whether inference to the best explanation (IBE) is a generically valid form of reasoning (especially for its applicability to the scientific realism debate). Attending to the specific roles that explanation plays in scientific methodology, especially the way it structures discourse in a discipline and coordinates exemplars, reveals the possibility of justifying explanatory power as an epistemic virtue in specific scientific cases, without reliance on general philosophical arguments based on IBE or underdetermination. This kind of argument naturally requires close attention to the historical development of a theory and its applications. Inflationary cosmology, I claim, offers just such a compelling, concrete example.
Inflation is a cosmological scenario that was originally proposed in the early 1980s by Alan Guth. It was widely accepted in the community immediately after its introduction, and remains a central pillar of the contemporary standard model of cosmology. Inflationary theory is based on the supposition that the very early universe underwent a brief period of accelerated and exponential spatial expansion. Proponents claim that the effect of inflation is to flatten the spatial geometry of the universe and make its contents more uniform. (One may usefully compare it to the inflation of a balloon, which decreases the curvature of the balloon's surface and smooths small irregularities.) This mechanism is thought to operate for a short period in the very early universe, giving rise to the conditions that eventuate in the present spatial flatness and uniformity, conditions which we infer today from observations of the cosmic microwave background (CMB) radiation. Proponents also claim that the old standard cosmological model, the well-known hot big bang (HBB) model, suffers from fine-tuning problems. Earman and Mosterin have emphasized that these fine-tuning problems are not problems concerning the HBB model's consistency or empirical adequacy, since the model is capable of explaining the present flatness and uniformity of the universe; rather the problems appear to raise concerns over the kind of explanation given by the model for certain physical features of the universe which are accessible to observation. In particular, only explanatorily-deficient special initial conditions can give rise to these presently-observed conditions within the context of the HBB model. Since uniformity and flatness are thought to be natural outcomes of inflation, the previous paradigm's fine-tuning problems are apparently solved by inflationary theory, thus leading to the claim that inflationary models represent real theoretical progress over the HBB model.
Although inflation was widely accepted (ostensibly on the basis of such fine-tuning arguments) during inflationary theory's early history, at present the best argument for inflationary theory is not that it (allegedly) solves these problems; instead it rests on the striking empirical confirmation in the 90s of quantum mechanical predictions developed out of the inflationary framework, specifically of a very precise spectrum of anisotropies of the CMB. If this latter, empirical argument is successful (it at least appears to be taken as such by most contemporary cosmologists), then inflationary theory should reasonably be considered an empirically successful theory whose predictive successes go beyond the HBB model, and therefore represent progress over it. Yet it is important to note that these predictions were unforeseen at the time of inflation's proposal and initial acceptance. How then is it, that a theory, seemingly unjustified on any of the commonly accepted epistemic grounds, should later find itself strikingly confirmed observationally? The case appears to be one of extraordinary luck, i.e. epistemic success achieved through a method no more effective than guessing. Yet supposing it so is quite implausible, for this degree of luck in the confirmation of central scientific theories would severely threaten any intuitive notion of scientific progress and rationality. The alternative to such skepticism is to consider that inflation's rapid and early acceptance among cosmologists was somehow epistemically justified prior to any observational support, and on grounds other than observational support or solving theoretical inadequacies in the Standard Model. Therefore the case of inflation shows us that a view of epistemic justification based solely on the simple and familiar criteria of empirical adequacy and theoretical consistency are seriously inadequate.
I claim that the epistemic justification of inflationary theory (before its observational confirmation) rests instead crucially on explanatory considerations, considerations which may be seen to arise from its approach to solving the HBB model's fine-tuning problems and explaining presently observed cosmological conditions. One might wonder, "How can solving such mere explanatory problems represent progress towards an empirically successful theory?" Insofar as scientific progress may be gauged by solving scientific problems (a la Kuhn or Laudan), one has, I claim, an explanatory story linking inflationary theory's putative success at solving the HBB model's fine-tuning problems with the later confirmation of its observational predictions. Roughly speaking, one might say that by solving the HBB model's conceptual problems, inflationary theory proves itself to be a progressive research program suitable for further development and empirical test. This viability depends on a certain kind of "meta-empirical" confirmation. Although, certainly, there is no guarantee that its predictions will be borne out, one's confidence in the theory is justified by its past problem-solving success. The viability of some such story of course depends however on whether inflation does in fact solve the HBB model's fine-tuning problems in some sense. Nevertheless, this argument sketch makes considerably more sense of the historical development of inflationary theory than an impoverished, simplistic view of empirical justification in science can.

Joseph Melia. Fundamentality, Composition and Modality Claims involving fundamentality, or the relatively fundamental, are common in both philosophy and physics. It is frequently claimed that our deepest physical theories, such as General Relativity and Quantum Field theory are more fundamental than theories of biology or sociology. But in many such cases, the notion of fundamentality can plausibly be understood in an epistemic or semantic sense -- as a relationship between theories. What is more contentious is whether and how claims about fundamentality can or should be understood in a metaphysical or ontological sense. Yet, in recent years, in both the philosophy of science, this concept is used to articulate and formulate certain key views.
For instance, in the philosophy of science, non-eliminative ontic structural realists distinguish their positions from realism simpliciter by saying that, although ordinary objects such as tables and chairs exist, their existence is, in some sense, derivative or secondary: at the fundamental level, all that exists is struc-

ture (Ladyman, 1998). In recent versions of the Many Worlds interpretation of Quantum Mechanics, the Many Worlds themselves are said to be emergent or derivative entities, not part of the theory's fundamental ontology (Wallace, 2014). The claim that it is symmetries, rather than particles, that are fundamental to nature is made by both physicists and philosophers (Heisenberg, 1975). And, in quantum information theory, Wheeler and others have argued that the correct lesson to be learned here is that information itself is ontologically fundamental and ontologically prior to matter (Wheeler, 1990).
Clearly, the notion of Fundamentality has come to play a pivotal role in the articulation and definition of a number of important positions in philosophy of physics and physics itself. Yet the concepts of fundamentality and ontological priority is frequently left undefined or sketchy by workers in the philosophy of science. Without a clear grasp on the relevant concepts, we cannot be said to have a clear grasp on the positions that use the concept in their articulation. An understanding and elucidation of Fundamentality is therefore a vital part of the articulation of these views.
In this paper, I examine competing modal conceptions of the relatively fundamental. I argue that, in recent years, the metaphysicians' conception of the fundamental has become divorced from the practices of physics: for those of us who wish to pursue philosophy along broadly Quinean, Naturalistic lines, this situation is unsatisfactory. Moreover, while many metaphysicians seem to have become comfortable with fine-grained and hyper-intensional notions of modality -- such as Finean essences -- I argue that certain Quinean worries have not been met satisfactorily: these worries make essences and similar notions ill suited for the naturalist.
I argue that the physicists' conception of a state space supports a conception of the modal which is, with some caveats, broadly Humean and combinatorial in spirit. I develop such an account and use it to show that a supervenience-based account of the fundamental can be developed which manages to avoid a number of the standard objections that have been used to discredit such attempts: that supervenience relations are not explanatory, that they merely signal a co-variation, that they are not fine-grained enough can all, I claim, be met: provided supervenience is understood in terms of the right modal concepts.
What supervenience alone, however, doesn't give us, is a completely satisfactory notion of priority: yet some such notion seems to be required of the ontologically fundamental. For this, I argue, we need some notion of composition, of the part-whole relation. The kind of composition we need may not be mereological, but it is the kind of composition that holds strictly between particulars -- tables, atoms, electrons: not properties, states of affairs or facts. Composition, coupled with the right definition of supervenience, can be used to give an account not only of fundamental particulars -- those objects that have no proper parts -- but can also give a (limited) account of the notion of a fundamental property. But the resulting account is limited only in the sense that it does not do everything that philosophers have wanted from the notion of a fundamental property (there is, for instance, no unique base level of fundamental properties) and will not serve all the purposes to which philosophers have wished to put fundamental properties, it is, I argue, sufficient to serve the purposes of mainstream science -- and that is sufficient for the naturalist.

Tushar Vasudev Menon. Taking up superspace- what would it take to be a realist about superspace? Modern supersymmetric theories present an interesting interpretative challenge. As a result of consistency conditions on the algebra of the supersymmetry (SUSY) generators, one is led to the idea that SUSY, although traditionally defined as a dynamical symmetry between bosons and fermions, could also be thought of as a spacetime symmetry in some extended spacetime, known as superspace. SUSY is, among other things, a crucial part of the string theoretic framework for a theory of quantum gravity. This talk focuses on what it would take to argue for an interpretation that favours the superspace formulation. After setting up the relevant terminology and distinctions, I introduce a stripped down toy model of a supersymmetric field theory and argue for a special case of a more general thesis--- that one needs some pre-existing philosophical commitment to favour one mathematical formulation over another. I then consider three extant arguments from the literature on the philosophy of spacetime as candidates for such a position in the context of supersymmetric theories.
Identifying an appropriate spacetime structure and material ontology for a dynamical theory in physics is not a trivial task. The standard way of doing so is by appeal to symmetry arguments in some guise or another, which are then used to impose restrictions on the models corresponding to a theory. The broad purpose of this paper is to consider the role that SUSY plays in constructing an appropriate spacetime for a supersymmetric dynamical theory. In practice, what this amounts to is an assessment of the viability of a realist stance on superspace.
Superspace is an extension of our ordinary four-dimensional spacetime to include (at least) four new dimensions, coordinatised by anticommuting elements of a Grassmann algebra. It is not immediately clear what a realist interpretation of superspace is. It does not amount to being a substantivalist about superspace because both substantivalists and relationalists about space/spacetime/superspace agree on certain elements of the mathematical setup, and it is these elements which I take the realist interpretation of superspace to question. The relationalist--substantivalist debate is therefore orthogonal to the `reality of superspace' debate as I construe it. The former debate is, at its core, concerned with what constitute the relata of the geometric relations that constitute part of a dynamical theory. The latter debate, on the other hand, questions the aforementioned assumption that both substantivalists and relationalists agree on--- that spacetime relata, if they are to be coordinatised at all, are best coordinatised as quadruples of real numbers.
A choice of coordinatisation proceeds in two steps. The first is to decide on a suitable mathematical space, the second is then to pick out privileged n-tuples of this space which are best adapted to the descriptions of spacetime entities of concern in a theory. The debate about the reality of superspace is concerned with how well we can motivate the claim that the best mathematical space to represent the ideology of our theory is a Grassmann algebra, rather than a commuting algebra R^4.
I therefore construe a realist interpretation of superspace as an interpretation under which the coordinates assigned to spacetime relata (whatever they may be) include ordinary real number coordinates (the commuting components of a Grassmann algebra) as well as anticommuting elements of a Grassmann algebra. My use of the term `ideology' is derived from Quine and refers to the properties and relations in our theory which we deem to be primitive.
This debate has implications for our view of the nature of spacetime as described by string theories. Any (string) theory which purports to describe our world needs to have the conceptual resources to describe the dynamics of fermions. The only way to incorporate fermions into a string theory is through supersymmetry. So any thesis regarding spacetime in a string theory that hopes to model our actual world must include reference to the role that supersymmetry plays. It has been claimed that the vibration of strings itself is a vibration in superspace. In this talk, I choose to focus only on supersymmetry, divorcing it from the context of string theory.
On the superspace formalism, SUSY can explicitly be thought of as a spacetime symmetry. But that is not enough- it can equally well be thought of as a dynamical symmetry on ordinary spacetime, according to the standard formulation. What we need is a further metaphysical or epistemological (or methodological) principle to give us independent reasons to prefer the spacetime formulation of SUSY over the dynamical one.
This is where a pre-existing philosophical commitment can be used to break the interpretative stalemate. After all, we cannot look at the mathematics to suggest which model we should interpret realistically. Earman's Principle, that one should match one's dynamical and spacetime symmetries exactly, is one option. The second option is to consider a view on which two situations related by a symmetry are the same situation differently described. This so-called symmetry-as-redescription interpretation (SRI) is a metaphysical position that, in some sense, underwrites Earman's methodological principle. The SRI can be cashed out in many ways- in this talk, I choose to use Saunders' appeal to Leibniz's Principles. A third option is to approach the supersymmetric generalisation of Minkowski space by analogy with the dynamical approach to ordinary Minkowski space, as advocated by Harvey Brown. This also underwrites Earman's Principle, but in a different way to the SRI.
This is not to suggest that these are the only three options at hand, it is merely to suggest a possible way forward. Several further problems need to be dealt with when applying this notion either to supersymmetric field theories or superstring theories. But what I present here is a first attempt at identifying some of what might go into an interpretation of a spacetime compatible with SUSY.

Thomas Moller-Nielsen. Invariance, Interpretation, and Motivation Take the 'Invariance Principle' to be the principle that only quantities that are invariant under the symmetries of our theories are physically real (cf. Saunders 2007). It is a doctrine with a distinguished pedigree: acclaimed theorists as diverse as the physicist Paul Dirac (1930, vii), the mathematician Hermann Weyl (1952, 132), and the philosopher Robert Nozick (2001, 82) were all apparent signatories during their respective lifetimes. Prima facie, however, it is something of a mystery as to how and why the principle is supposed to work. Nevertheless, there appear to be at least some uncontroversial cases where it---or something very close to it---does work.
One such example can be found in Newtonian Gravitation Theory (NGT), i.e., the theory comprising Newton's three laws, plus his inverse gravitational square law, governing the behaviour of point particles in Newtonian spacetime. As is well known, this theory is Galilean invariant. This implies, among other things, that if one takes any solution to NGT and 'boosts' it---that is, uniformly alters the absolute velocity of each point particle by the same amount throughout its history---one will invariably get back a solution to NGT. Boosts, in other words, are a *symmetry* of NGT: they are transformations that invariably map solutions of the theory to solutions.
Which quantity varies under this particular symmetry? The answer is obvious: absolute velocity. Thus, according to the Invariance Principle, we should conclude that absolute velocity is not a genuine physical quantity. Conversely, which quantities are invariant under this particular symmetry? Again, the answer is obvious: relative (inter-particle) distance and velocity, temporal intervals, and absolute acceleration. Thus, according to the Invariance Principle, we should conclude that NGT's boost symmetry does not threaten these quantities' status as genuinely physical.
As it turns out, one can successfully purge Newtonian theory of the spacetime structure required to make absolute velocity a physically meaningful quantity. More specifically, one can move to 'Galilean spacetime'. (Sometimes also called 'Neo-Newtonian spacetime'.) Here, the Newtonian posit of persisting points of absolute space---persisting points which, crucially, allow for the notion of absolute velocity to be physically meaningful---is done away with, but an affine structure is nevertheless preserved, which defines the 'straight' or force-free (inertial) paths through spacetime. Absolute velocity is therefore not a physically meaningful quantity in Galilean spacetime, as it is in Newtonian spacetime. Nevertheless, all other Newtonian notions, including the notion of absolute acceleration, remain well-defined in Galilean spacetime. To the extent that one opts for Galilean over Newtonian spacetime, then, one has excised an ostensibly odious piece of theoretical structure from NGT.
Three important caveats are worth noting, however. First, and most obviously, none of this is to say that Newtonian theory set in Galilean spacetime is therefore the true and complete theory of the world. (It isn't.) Second, nor is this to say that by moving to Galilean spacetime one has thereby purged Newtonian theory of all its 'variant' structure. (One hasn't. The symmetry group of Newtonian theory is actually wider than the Galilean group: it has additional symmetries; cf. Knox 2014.) Third, nor is this even to say that the invariant quantities one ends up with following such an application of the Invariance Principle will invariably be preserved in future theories. (For instance, there is no notion of 'relative spatial distance' simpliciter in special relativity.) Given all of these caveats, however, one might well ask: What good is the Invariance Principle, exactly? What purpose, in particular, does it serve?
As I see it---and, I take it, as many other contemporary theorists also see it---the purpose of the Invariance Principle is essentially *comparative*. That is, it is simply supposed to lead you to a *better theory*---or a better 'interpretation', or characterisation, of the same theory---than the one you started with. To take the case at hand: Newtonian theory as set in Galilean spacetime is a better theory than Newtonian theory as set in Newtonian spacetime. It is a theory which possesses all of the theoretical virtues of its rival, but lacks any apparent ontological commitment to the unwanted variant quantity in question.
In summary, the Galilean invariance of NGT, in conjunction with the Invariance Principle, is supposed to indicate that neither absolute velocity nor any corresponding persisting points of absolute space are genuinely real. Now to lay my cards on the table: I actually think that something *very close* to this general kind of inference---that is, from the variance of a quantity under symmetries to that quantity's nonreality---is legitimate. The devil, however, is in the details. In particular, I don't believe that the *mere* Galilean invariance of NGT is enough to establish absolute velocity's nonreality. And in general, I don't believe that the *mere* variance of a quantity under symmetries is enough to establish that quantity's nonreality. These beliefs, as far as I can determine, put me in the

minority camp in the contemporary philosophical literature on symmetries. Nevertheless, I think they are correct beliefs---and they are the ones that I will attempt to argue for in the remainder of this talk.
The structure of the rest of the talk is as follows. First, I distinguish between two importantly distinct ways of thinking about symmetries: what I call the 'interpretational' and 'motivational' conceptions of symmetries. (In brief, according to the (orthodox) interpretational view, if two solutions are related by a symmetry, it is always legitimate to *interpret* them as representing the same physical state of affairs. According to the (unorthodox) motivational view, however, we are only ever *motivated to regard* symmetry-related solutions as physically equivalent.) I then move on to extol the benefits of the motivational view, and distinguish it from the related but, I claim, erroneous view that symmetries invariably prompt a mathematical reformulation of the relevant theory. I finish the talk by launching a more direct series of attacks against the interpretational view, which hinge largely on the claim that it is incompatible with a form of scientific realism worthy of the name.

Andrea Oldofredi. Particles creation and annihilation: a Bohmian Approach Physics has always been concerned with questions regarding what are the ultimate constituents of matter and how they behave and interact. The Standard Model (SM) of particle physics is an answer to these questions, and nowadays is the most successful physical theory at our disposal. This model explains the fundamental structure of matter in terms of elementary fermions interacting through bosonic fields and comprehends three of the four fundamental forces in Nature: the electromagnetic, the weak and the strong interactions; only gravitational effects are not taken into account. Furthermore, its predictions have been corroborated with an extreme degree of accuracy, and recently remarkable experimental evidence for the existence of the last ingredient of the SM, the Higgs boson, have been obtained.
The SM is a Quantum Field Theory (QFT), in the sense that QFT is the mathematical framework in which SM is written. Thus, it provides an ontology in terms of fields, and it is by construction a unification of the axioms of Quantum Mechanics (QM) and Special Relativity (SR). (It is worth noting that in experimental situations, in order to produce new particles from collisions, energies are needed to be at least as great as the rest masses of the produced particles, thus relativistic requirements must be necessarily taken into account. Moreover, SM predicts the existence of antiparticles which come from the negative solutions of the Dirac equation as consequence of the relativistic relation E=±√(p^2 c^2+m^2 c^4) present in it. These are only two of the several reasons according to which it is not possible to dismiss relativity in QFT.)
Despite of these significant triumphs, this theory inherits several conceptual problems that plague the standard interpretation of quantum mechanics, such as the measurement problem or the role of the operators and of measurements. Thus, mathematically ill-defined notions appear even in the fundamental structure of the SM.
However, among the foundations of quantum physics, there exist models with a clear ontology, e.g. Bohmian Mechanics (BM) or the spontaneous collapse theories (GRWm, GRWf, rGWRf), in which such notions do not find any room within the derivations admissible from their axiomatic apparatus. The primary aim of this talk is to present the common structure of these theories underlying the crucial role that a sharp ontology plays in order to obtain successful explanations of physical phenomena.
In second place, two models of Bohmian QFT will be presented as serious alternatives to the standard formulation of QFT, in order to recover the physical content of the SM. Though standard QFT is generally defined as the combination of the axioms of quantum mechanics and special relativity, there exists a class of non-relativistic models which are generalizations of Bohmian Mechanics to the phenomena of particles creation and annihilation reproducing the statistics of QFT experiments. In this talk, I will presents two models which share a particle ontology, being insensible to the conclusions of several no-go theorems which exclude the possibility of a proper particle theory in the context of QFT. (they involve specific relativistic constraints which are violated in BM) These are the Dirac sea approach and the Bell-type QFT. The former postulates an ontology of a finite and fixed number of fermions, which are defined as structureless particles with a specified position at every time t. Within this model particles are never created or destroyed. The dynamics is completely deterministic and comprehends the usual Schrödinger equation for the motion of the wave function and a guiding equation describing particles' trajectories. Here bosonic degrees of freedom are not part of the fundamental ontology. The latter provides an ontology made of fermions and bosons both considered as elementary particles (with positions always defined). The dynamics for the configuration of particles is stochastic: here trajectories can begin and end, therefore, random jump processes from a given configuration to another are inserted within a Bohmian-like guiding equation. These jumps are related to creation and annihilation of particles. Though this model does not provide a deterministic law of motion, it reproduces the statistics of the standard model considering equivariant Markov processes by construction.
Even though these models are not relativistic, they gain Lorentz invariant predictions, being experimentally indistinguishable with respect to a genuinely relativistic theory.
These models show that it is mathematically possible to postulate a particle ontology even in QM and QFT, providing an image of the world approximately similar to that of classical physics. These results are achieved specifying a primitive ontology (determination of the fundamental entities the theory is about) and a set of dynamical variables which constraints the motion of the primitive variables. This strategy follows the methodology introduced in the mid-Seventies by the physicist John S. Bell. In several papers he explained how to construct a rigorous physical theory from a sharp metaphysics. This methodology divides the mathematical structure of a given theory in two parts: structures with a direct physical meaning and dynamical structures. The former ones are the formal counterparts of real physical objects postulated as primitive concepts according to a specific theory. Since they are always localized in space and time they are called local beables (for instance, in BM the local beables are particles' positions). These primitive variables cannot be defined in terms of other more basic notions and the explanation of every physical phenomenon is based on them. The dynamical structures are used to implement equations of motion for the former ones: they tell how these move in space and time via the specification of parameters such as mass, charge, energy, wave functions, etc. Considered together these two structures define the "architecture" of a physical theory.
In conclusion, the substantial aim of the talk is to underline the how a clear metaphysics at the fundamental level of construction of physical theories could be extremely useful in order to avoid the severe conceptual problems that plague the standard version of QM and SM (or more generally QFT) and to achieve rigorous physical theories. The Bohmian QFT models here considered are interesting examples of how this goal could be obtained.

Viorel Pâslaru. Mechanisms, Predictions and Explanations The new mechanistic philosophy defined itself relative to the explanatory practice of biologists and in contrast to the DN model. Accordingly, the main function of descriptions of mechanism is to provide explanations. Due to this focus on explanation, new mechanistic philosophers have given little attention to the mechanistic basis of prediction despite the central role of prediction in the sciences. In this paper, I argue for a way of correcting this deficiency in the study of mechanisms.
In many sciences, such as ecology, which is my area of focus, the role of descriptions of mechanisms in formulating predictions is on a par with, but often more important than their role in explanations. A major goal of describing mechanisms is to formulate predictions (Tilman 1987, [1990] 2003, Johnson et al. 2013, Mouquet et al. 2015). Descriptions of mechanisms for the purpose of prediction are not different from descriptions of mechanisms for the goal of explanations. Even when ecologists do not find causal-mechanistic explanations useful and argue for a predictive ecology, they nevertheless base predictions on mechanistic descriptions of phenomena under scrutiny (Peters 1991, Shipley and Peters 1991). Yet philosophers of mechanisms have not examined this role of describing mechanisms. Glennan (2002, [2008] 2014) does not discuss the link between mechanistic explanations and predictions. Machamer, Darden and Craver (2000), and Craver and Darden (2013) affirm that mechanism schemata and descriptions of particular mechanisms are used to predict phenomena; however, they do not elaborate on this topic. Bechtel and Abrahamsen (2005) briefly examine the relationship between prediction and mechanisms in the context of challenges associated with tests of mechanisms, but they do not scrutinize the problem. Williamson and Illari's (2012) consensus account of mechanisms acknowledges their role in prediction along with explanation, but it does not shed light on the specifics of this role. Even philosophers who examine prediction in ecology (Elliott-Graves 2015) have not considered the role of mechanism descriptions in making predictions possible.
Against this background, I argue for using Hempel's symmetry thesis as a criterion of adequate explanations and show that correct predictions validate mechanistic explanations formulated based on a notion of mechanism different from those defended in the new mechanistic philosophy. To accomplish this objective, I outline first a conception of mechanisms that by contrast to established accounts (1) de-emphasizes organization as a stand-alone factor in determining the operation of mechanisms, but underscores the role of properties in deciding the organization of mechanisms and the kind of activities that entities engage in; (2) acknowledges the role of causation by omission in determining ecological phenomena; and (3) asserts that insensitive and invariant causal networks are the basal structure of mechanisms. I describe then how mechanisms conceived in this way underlie predictions. After that, I distinguish two types of explanation in Hempel's DN model. Hempel and Oppenheim (1948) drew this conclusion from the thesis of structural identity: "It may be said, therefore, that an explanation of a particular event is not fully adequate (my emphasis –V.P.) unless its explanans, if taken account of in time, could have served as a basis for predicting the event in question" (138). I interpret this condition of full adequacy of explanations as outlining two categories of explanations. Fully adequate explanations are explanations that scientists should strive to offer, i.e., explanations that can be used to formulate predictions. As such, the condition of full adequacy formulates an ideal for scientific explanations. By contrast, explanations that cannot be used for predictions form the category of partially adequate explanations that fall short of the ideal, even if they satisfy the conditions of logical and empirical adequacy and answer explanation-seeking questions. In light of the foregoing, I argue that the condition of full adequacy articulates a requirement for improving explanations of the second category. Fully adequate explanations describe the functioning of mechanisms in a wide variety of conditions and it is this features that allows them to operate as predictions. Finally, I show that my approach does justice to how ecologists view the relationship between mechanistic explanations and predictions (Peters 1991). In addition to satisfaction of constitutive relevance exhibited by means of top-down and bottom-up experimental strategies (Bechtel and Richardson 1993, Craver 2007), formulation of correct predictions is an important way to ensure that explanations are adequate.

J. Brian Pitts. Unconceived Alternatives, Space-time and Particle Physics Einstein developed General Relativity prematurely. The mainstream development of relativistic quantum mechanics involved a systematic exploration by many over decades. All wave equations that are at least special relativistic appear in Eugene Wigner's taxonomy in terms of 'rest mass' (an inverse length m) and 'spin' (intrinsic angular momentum, related to scalar vs. vector vs. matrix potential). A massive particle/field has waves with a frequency-dependent speed (not c) and a point source solution exp(-m*r)/r.
Particle physics has a strong, largely positive impact on the rationality of gravitational physics in comparison with Einstein's somewhat mysterious development of General Relativity (noticed by Norton and Renn among others). In 1939 Pauli and Fierz recognized linearized General Relativity as mass 0, spin 2. Could one derive the complicated nonlinearities as well? This project succeeded in the 1950s-70 due to Kraichnan, Gupta, Feynman, Deser and others. From the early 1970s until the 2000s, particle physics offered an eliminative induction for General Relativity: basically, write the most general relativistic wave equation with light bending (tensor rather than scalar), require stability and so exclude negative energies, infer 'improper' conservation laws in the massless case, and arrive at Einstein's equations using (in effect) Noether's converse "Hilbertian assertion" from improper conservation laws to general covariance. A key result from the 1970s (until recent Kuhn loss!) refuted massive spin 2. Thus gravity, if it bent light, could hardly fail to satisfy Einstein's equations.
While particle physics excludes many competitors to Einstein's theory (such as Rosen's 1970s bimetric theories and Einstein's 1913-15 Entwurf) due to unstable negative energies, it also suggests perhaps the most serious competition. In the 19th century Seeliger and Neumann modified gravity on long distances with a new length scale. From the 1920s-30s (de Broglie's massive photons, the Klein-Gordon equation, Yukawa's work), this inverse length acquired a physical meaning: the graviton rest mass. A sufficiently small graviton mass should be empirically indistinguishable from 0, giving permanent underdetermination from approximate but arbitrarily close empirical equivalence, Seeliger's theme that would not have surprised Duhem. That expectation is fulfilled for many theories. For spin 2 gravity, devils appeared in the details in the 1970s. But they were

exorcised by new or rediscovered results in 2010-11 by de Rham, Gabadadze, and Tolley, and by Vainshtein and others in the 2000s, evading old folk impossibility theorems.

Massive spin 2 gravity is perhaps the most serious competition to General Relativity considering philosophical expected utility involving prior plausibility, presumptive empirical adequacy, and philosophical payoff. It shows how Einstein's principles (generalized relativity, general covariance, and equivalence) admit loose readings that are plausible but not very interesting, and exact readings that are interesting but conjectural, as Freund, Maheshwari and Schonberg noted in 1968. Mass terms shrink symmetry groups; de Broglie-Proca massive spin 1 electromagnetism has a physically meaningful electromagnetic 4-potential Aμ satisfying deterministic field equations, whereas Maxwell's theory (massless spin 1) has gauge freedom. Spin 2 gravity prima facie is analogous.

Massive gravity should have moderated the revolutionary zeal of some proponents of Einstein's theory, not least logical empiricist Moritz Schlick's critique of Kant. Schlick, leader of the Vienna Circle, transformed philosophy based on the supposed empirical refutation of synthetic a priori knowledge. Schlick reasoned that if geometry demonstrably was not a priori, then nothing was. This project was scientifically avoidable at the time; one only needed to do to Einstein's theory what Neumann, Seeliger, and Einstein (1917!) did to Newton's, to make space-time physics Kant-friendly. Massive gravity involves two geometries, one flat (hence a priori)---an instance of a more general geometry invented by Levi-Civita no later than 1923 but with hints in Lobachevsky in the 1830s. If massive spin 2 gravity had been invented in 1917 or 1923 instead of 1939, and especially if it had been noticed by philosophers, it would have blocked Schlick's argument overthrowing Kant's synthetic a priori knowledge. A call to overthrow Kant might have arisen instead when massive spin 2 gravity collapsed---which happened in 1972 (if even then, Maheshwari having partly anticipated de Rham et al. but making no impact). The recent collapse of 1970s arguments against massive spin 2 gravity, surprisingly, partly reopens the door to Kant. Massive spin 2 gravity was neglected early on partly due to Einstein's 1917 false analogy between (in effect) a graviton mass and his new cosmological constant Λ---first detected in the 1940s by Heckmann, hardly discussed prior to the 1960s, and still tempting to historians. In the long run particle physics should provide either a better argument for Einstein's equations (eliminative induction) or a serious rival theory.

Recent General Relativity historiography by Renn, Sauer, Janssen et al. shows that Einstein found his field equations partly by a physical strategy including the Newtonian limit, the electromagnetic analogy, and energy conservation. Such themes are similar to those later used by particle physicists. How do Einstein's physical strategy and the particle physics derivations compare? Given that Lagrange, Hamilton and Jacobi linked rigid translation symmetries and conservation laws in the first half of the 19th century, did Einstein? How did his work relate to emerging knowledge (1911-14) of the canonical energy-momentum tensor and its rigid translation-induced conservation in Herglotz, Mie and Born? He diligently sought conservation laws using linear coordinate transformations and the gravitational field equations alone, but seems not to have recognized that conservation was guaranteed using merely the uniformity of nature (rigid translation invariance, which he tacitly assumed) and the field equations for gravity and matter.

What was really wrong with Einstein's Entwurf theory? In 1915 he retrospectively faulted it for not admitting rotating coordinates and not getting Mercury's perihelion right. Neither aspect is fatal, but its having 3 negative-energy field degrees of freedom (failing a 1920s-30s particle physics stability test with antecedents in Lagrange and Dirichlet) would have doomed it regardless.

Particle physics thus can be useful in the study of gravity both in assessing the growth of objective knowledge and in suggesting novel lines of historical inquiry for the historiography of General Relativity. The history of General Relativity assists particle physics via Noether's converse results.

Carina Prunkl. Is there a thermodynamical cost associated with some interpretations of quantum theory? For nearly a century now we have puzzled over how to interpret quantum theory, unable to decide between a variety of promising candidates. Recently, however, an argument has been brought forward by Cabello et al. (2015), claiming to restrict the class of admissible interpretations radically. On the basis of information theoretic considerations, the authors assert that there exists an empirically testable difference between two broad classes of interpretations. According to their argument, popular approaches like Bohmian mechanics, Many Worlds and GRW are left untenable as in such approaches non-orthogonal, successive measurements on an individual system are seemingly associated with an implausible heat generation. I will argue that their argument is mistaken by showing that it is based on a profound misconception of the underlying information theoretic notions - an endemic phenomenon in the physics literature. In clarifying the relationship between thermodynamics and information, I thereby aim at shedding some light on an often encountered misapprehension and to restore good sense.

Cabello et al. divide quantum interpretations into two broad classes, based on their varying approach to probabilities: Type-I interpretations, such as Bohmian Mechanics, Everett and GRW, take probabilities to be determined by intrinsic properties of the quantum system. Type-II interpretations, such as Rovelli's Relative-State Formulation and QBism, instead take probabilities as being ``relational experiences between an observer and the world''. The distinction is contestable, as some of the Type-I interpretations listed above cannot unambiguously be placed in the Type-I camp, but we will put this issue aside. The authors then consider successive measurements on an individual quantum system and make three assumptions: (i) which measurement is performed on the system is decided independently of the state of the system, (ii) the quantum system only has a limited memory and (iii) Landauer's principle is valid. Prima facie, these assumptions seem reasonable, but as I shall argue the meaning of (ii) is insufficiently clear, and Cabello et al. radically mis-apply (iii).

The simple version of their argument is quickly explained: we consider an individual two-dimensional quantum system which is subject to successive random measurements in either the x- or z-basis. Modelled as a stochastic input-output process (Crutchfield and Young, 1989), the Shannon information of the set of what they term `causal states', in this example the set of all possible current quantum states, is then said to quantify the information stored by the system in order to optimally predict the future state. Cabello et al. argue as follows: since the system can only store a limited amount of information (assumption ii), it needs to ``overwrite'' or rather ``erase'' information in order to have room for the intrinsic properties that determine its future behaviour. They suggest that this quantity is given by the conditonal entropy of the past causal state, given the current causal state and the choice of measurement basis. As a result, they maintain that due to Landauer's principle, such successive non-orthogonal measurements must result in the dissipation of the corresponding amount of heat into the environment. This would seem a reductio ad absurdum of Type-I interpretations, or in any case, it could be checked directly by experiment.

Groundbreaking as this may seem at first sight, the stated results cannot withstand a deeper analysis. In particular, as I will explain, we can construct a direct counterexample to their argument by considering the setup applied to Spekkens' toy model, which provides a Type-I interpretation of this experimental scenario. Non-orthogonal measurements can be associated with the horizontal or vertical insertion of a partition into a box containing a single particle, in contact with a heat bath. In this model it is immediately clear there need be no heat exchange with the environment.

This leads us to the main point of discussion: the relation between information and thermodynamics - a relation that is frequently associated with Landauer's principle. The principle is often taken to be the statement that the implementation of a logically irreversible operation is accompanied by the dissipation of heat. Being only true in some special cases however, this particular phrasing of Landauer's principle needs to be tested against a more general version that takes into account the relation between logical and physical states of the system in question (Maroney, 2009). A concrete counterexample to the above phrasing is an operation that randomises the logical state of a bit: it is logically irreversible but can be implemented thermodynamically reversible and without any associated heat cost. Furthermore, the thermodynamically relevant quantity for Landauer's Principle is the von Neumann entropy, as opposed to the Shannon entropy, to which there is no general one-to-one correspondence. This also undermines the idea that information, as quantified by the Shannon entropy, is some kind of substance, whose change inevitably has thermodynamic consequences.

A second issue is that Cabello et al. are effectively insisting that there is a thermodynamic relevance of epistemic uncertainties about the past. There are of course cases in which epistemic uncertainties track the underlying physical state of the system and where they are of importance - but not here. Whatever intrinsic features determine the probabilities of the system at a time, they are not affected by the occurrence or non-occurrence of an event at a later time. Uncertainty about the past hence cannot be taken to determine the intrinsic features of the system in the past and, in particular, uncertainty about the past does not have any current thermodynamic relevance. For example, I don't heat up just because I forgot what I had for breakfast this morning, but as I shall explain, it is on inferences of this sort that Cabello et al.'s argument relies.

The disconcerning claim that successive non-orthogonal measurements should be associated with a heat cost in Type-I interpretations can hence considered to be untenable. In showing this I hope I have not only clarified the application of Landauer's principle in the quantum context, but furthermore elucidated how and when information is of any thermodynamic relevance, hopefully preventing further misconduct.

Mantas Radzvilas. Preference Reconstructions and the Inner Best-Response Reasoner: A Critical Evaluation of the Explanatory Relevance of Payoff Transformation Models Experimental results suggest that people often choose strategies that do not constitute a Nash equilibrium, and report that they expect the outcome of the joint actions of the players to be a strategy profile that is not a Nash equilibrium of the experimental game. This pattern of "cooperative" behavior is observable in most experiments involving "social dilemmas" – games where individuals' personal incentives and social optimality diverge, such as Prisoner's Dilemma, Chicken game, and Public Goods game (for extensive overview, see Gächter and Thöni 2007).

In this paper, I focus on two approaches within the rational choice tradition that purport to explain such experimental results.

The received approach of behavioural economics is the payoff transformation approach. It is a modeling strategy which rests on assumption that the monetary payoffs of experimental games do not capture all the relevant motivations of players, meaning that people may be playing a different game where cooperative outcome is rational (that is, consistent with the best-response reasoning assumption). The explanatory strategy of this type of behavioural models is to "reconstruct" the utility functions of players, and show that cooperative outcome is a Nash equilibrium of the transformed game. Some of the payoff transformation suggest that people have stable pro-social preferences (such as Fehr and Schmidt's (1999) theory of inequity aversion), while other models suggest that people have conditional preference to act cooperatively in situations where they believe that other people are conditional cooperators as well (such as Bicchieri's (2006) theory of social norms).

An alternative modeling strategy is agency transformation approach. It is based on assumption that certain structural properties of games may trigger a shift in player's mode of reasoning from individualistic best-response reasoning to reasoning as a member of a group. When an individual reasons as a member of a group, s/he identifies the feasible strategy profile which is optimal for the group, and then plays his or her part in realizing it. Two of the better-known agency transformation theories are the team reasoning theory introduced by Sugden (1993, 2000, 2003) and Bacharach (2006), and Misyak and Chater's (2014) theory of virtual bargaining. The explanatory strategy of this type of models is to show that the experimental game in question has the structural properties that may actually trigger a shift in players' mode of reasoning, and that the outcome of the experimental game, given the criterion of group optimality offered in the theory, would be identified as the optimal outcome by the players who reason as members of the group.

The payoff transformation approach is preferred to the agency transformation on the grounds that payoff transformation models explain the experimental data while retaining the best-response reasoning (expected utility maximization in strategic context) assumption – one of the cornerstones of normative game theory.

In this paper, I consider two possible interpretations of the explanatory scope of decision-theoretic models. First, decision-theoretic models can be interpreted as approximations to the observed choices: people choose their strategies as if they were expected utility maximizers with pro-social preferences, or as if they were aiming to maximally advance the interests of the group. Second, these models can be interpreted as approximations to the actual process of reasoning by which people arrive at their choices.

I argue that if both types of models are treated as approximations to observed choices (the received interpretation in behavioural economics), then there is no compelling reason for treating the payoff transformation models as more explanatory relevant than agency transformation models.

I claim that payoff transformation models do not yield theoretical predictions regarding observable choices. Payoff transformation theories show that pro-social preferences transform the original mixed motive games into coordination games where cooperative outcome is a Pareto optimal Nash equilibrium of the transformed game. But the fact that a cooperative outcome is one of the Nash equilibria of the game does not imply that players will choose cooperative strategies, nor does it imply that players will end up playing a Nash equilibrium of the transformed game. The standard epistemic model of game theory does not explain how rational players form beliefs about each other's non-strictly dominated (rationalizable) strategies, and so does not yield unique theoretical predictions in games with multiple Nash equilibria (this result is due to Aumann and Brandenburger (1995)).
For comparative purposes, I use Sugden's (2003) theory of mutually advantageous play as an example of an agency transformation theory that can be operationalized to yield unique theoretical predictions in both mixed motive and coordination experimental games. I show that these predictions fit well with the available data from a wide range of experimental games, and argue that agency transformation models can better account for the observed dynamics of choice behavior in experimental games than payoff transformation models.
Finally, I consider the argument that payoff transformation models are more explanatorily relevant due to the fact that best-response reasoning model is the best available approximation to the actual process of reasoning by which players arrive at their choices. I find this argument problematic for three reasons. First, there is no empirical evidence that best-response reasoning is a better approximation to the actual process of reasoning than alternative models. Second, even if best-response reasoning model were shown a better approximation to actual process of reasoning than group-directed reasoning models, it is not clear why observed deviations from the theoretical predictions of the standard game theoretic model should be treated as expected utility maximizing choices in transformed games, rather than mistakes made by less than perfect best-response deliberators, or, in games with multiple rationalizable strategies, as expected utility maximizing choices made in light of incorrect beliefs about opponents' strategy choices.
Third, I argue that the dual-system theory, which is used as theoretical justification of the psychological interpretation of expected utility theory, does not warrant an assumption that the mental process of strategically sophisticated system 2 will generate the strategic mode of reasoning consistent with best-response reasoning, nor does it warrant an assumption that system 2 can operate independently from non-sophisticated system 1.

Katie Robertson. The Justification of Coarse-graining in Statistical Mechanics Both classical and quantum physics face the challenge of reconciling the time-asymmetry of macroscopic processes with the time-symmetry of the microscopic laws in physics. In statistical mechanics, answers to the challenge go back to Boltzmann's H-theorem and the derivation of the Boltzmann equation. In recent decades, there is a well-established framework (applying equally to classical and quantum cases), which I dub the Zwanzig-Zeh-Wallace (ZZW) framework that shows how the irreversible equations of statistical mechanics can be constructed from the underlying reversible microdynamics. These irreversible equations describe the temporal asymmetry associated with thermodynamic processes such as gases spontaneously expanding.
Yet this framework relies on 'generalized coarse-graining': throwing away some of the information describing the system (via a probability distribution) by using a projection operator. Historically originating with Gibbs, this procedure of coarse-graining has been repeatedly criticized in the literature: e.g. Redhead describes it as "one of the most deceitful artifices I have ever come across in theoretical physics".
In this paper, I will consider three objections and argue that these objections depend on a mistaken justification of coarse-graining. I then give alternative justification and relate it to emergence.
(1): The coarse-graining approach is not empirically adequate, in light of Hahn's spin echo experiment. I argue that the spin-echo case is not one where coarse-grained methods are expected to apply by considering the conditions under which the coarse-grained dynamics of statistical mechanics gives the same 'relevant' probability distribution to describe the system, as the microdynamics (a meshing condition which I call, following Wallace, `forwards-compatibility').
The second objection is based on the idea that coarse-graining distorts the description of the system and furthermore, the asymmetry solely arises from the repeated coarse-graining in the higher-level irreversible dynamics. Thus, (2): the coarse-grained asymmetry is illusory.
The third objection in the literature is that the coarse-grained quantities differ from other putative physical quantities such as energy and mass. (3): The worry is that the asymmetry is anthropocentric or subjective because the choice of coarse-graining depends on us.
Given these concerns and objections, coarse-graining seems in need of a justification. This task can be split into two:
(Choice) What is the justification for the choice of coarse-graining map?
(At all) Why is it legitimate to coarse-grain at all?
I argue that behind (2) the illusory and (3) the anthropocentric objections is a common - but mistaken - justification, which I call the measurement imprecision justification. According to this justification, our limited observational capacities imply that we cannot distinguish between the fine- and coarse-grained (probability distribution) descriptions of the systems. However, I show that this justification is neither necessary nor sufficient for answering (Choice) and (At all).
Instead I offer an alternative justification. Statistical dynamics are not (only) motivated by the calculational intractability of solving 10^23 equations (in the case of 1 mole of gas). Rather, there are definite macroscopic regularities (such as: 'gases expand' or 'entropy increases') that would be lost amongst a morass of detail at the microdynamical level, roughly speaking, even if we could solve 10^23 equations. Throwing away these details by coarse-graining allows us to describe and explain these higher-level patterns. In particular, coarse-graining allows the construction of irreversible equations that give us quantitative facts about the macroscopic regularities such as relaxation times and transport coefficients. Thus, coarse-graining abstracts to a higher-level of description (and this answers why we can coarse-grain at all).
Yet we don't just want to abstract to higher-level, we want a theory of the goings-on at this level. Abstracting to the centre of mass of all philosophers of science is unhelpful unless we can discuss the motion of this centre of mass without merely calculating the motion of each philosopher and then re-averaging. If we cannot say anything about what is going on at the higher-level of description without constantly referring back to the lower-level details, then the higher-level description will not be very useful. Thus, to have a theory of the higher-level, we need autonomous dynamics. In the ZZW framework, autonomy is one of the conditions required when constructing irreversible equations. But of course, not every coarse-graining map leads to autonomous dynamics. Thus, the choice of coarse-graining is determined by whether it leads to autonomous dynamics (and this answers the Choice justificatory question).
Having provided an alternative justification of coarse-graining, I can now answer objections (2) and (3) to coarse-graining. In reply to the illusion objection, coarse-graining is not an idealisation but an abstraction. Information is omitted rather than false assumptions added. Furthermore, the ZZW framework shows that - provided the system satisfies the meshing condition forwards-compatibility - the asymmetry is robust; it doesn't arise solely due to repeatedly coarse-graining. Hence, the coarse-grained asymmetry is not illusory.
In reply to the anthropocentric objection, this alternative justification (and more generally, the ZZW framework) shows that we do not choose the coarse-graining (and thereby taint the coarse-grained quantities with anthropocentricism). Rather the coarse-graining is justified (and useful) if it leads to an autonomous dynamics.
Finally, were coarse-grained best justified by the measurement imprecision justification, the asymmetry in statistical mechanics would be revealed to be subjective. Instead I draw a different moral: the coarse-grained asymmetry is weakly emergent. This conclusion is in direct opposition to Prigogine and Stengers who claim that "Irreversibility is either true on all levels or on none: it cannot emerge as if out of nothing, on going from one level to another". Whilst the dynamics at the lower-level of description is reversible, the dynamics at the higher-level of description is irreversible. True, this emergent irreversibility does not arise "as if out of nothing". Time-asymmetric assumptions are required in the ZZW framework. But this is to be expected; if no asymmetry is put in, then we cannot expect asymmetry out. In the ZZW framework, the added asymmetric ingredient is an initial condition - a particular form of the (controversial) Past Hypothesis. Lastly, I argue that my conclusion that the time-asymmetry is weakly emergent allows me to dismiss one of the worries about the Past Hypothesis.

Fiora Salis. Fictionalism about theoretical models and scientific representation Modern science crucially relies on idealized descriptions of hypothetical systems – or model systems – that scientists use to represent certain parts or aspects of the world – or target systems. A hypothetical system is chosen as the object of study because it is less complex than its target and because by using the model system to represent its target we can learn about the latter. So, the practice of modelling involves two steps. First, modellers prepare a model description of a hypothetical system as the object of study. Second, they use the hypothetical system as a representation of a real system in order to learn about the latter. This prompts the question of how models represent. Most philosophers of science phrase this problem in terms of necessary and sufficient conditions (one paradigmatic exception is Suàrez' (2003) deflationary approach). What they ask is how to fill the blank in the scheme 'M is a scientific representation of T if and only if ____', where 'M' stands for 'model' and 'T' stands for 'target'.
In this paper I will develop a new approach to the issue of models as representation that draws on Walton's (1990) aesthetic notion of fiction and the key notions of denotation and propositional imagination. To this aim I will individuate two main conditions that any model has to satisfy to be a vehicle of scientific representation. I will spell out Walton's aesthetic notion of fiction. I will assess current fictionalist accounts of models with respect to these conditions. I will spell out a novel account of models as representation in terms of the two key notions of denotation and propositional imagination. Finally, I will test the account on two case studies involving a theoretical model and a material model.
Something is a scientific model if and only if it satisfies two main conditions. The aboutness condition is that in virtue of which a model represents a certain physical system. Frigg (2006) refers to this as the enigma of representation, and Hughes claims that '[t]he characteristics – perhaps the only characteristic – that all theoretical models have in common is that they provide representations of parts of the world' (1997, 325). The epistemic condition is that according to which models enable scientists to formulate hypotheses and make predictions about physical systems. Suárez (2004) claims that models allow to draw inferences about physical systems, and Swoyer suggests that learning with a model involves a special cognitive function – or surrogative reasoning – that is 'directly about a representation in order to draw conclusions about the things that it represents' (1991, 449).
An early precursor of the analogy between models and fiction can be identified with Vaihinger's (1911/1924) emphasis on the importance of fictions for scientific reasoning. The analogy has been revived in different ways in recent literature on scientific explanations (e.g. Bokulich 2009, 2012), scientific models (e.g. Cartwright 1983, 1999; Morgan 2001, 2004), scientific epistemic practices (e.g. Elgin 1996), and more. Over the past decade, however, a more robust analogy has been explored by drawing on contemporary theories of fiction from aesthetics. Within this context the term 'fiction' applies to works of imaginative narration such as Flaubert's Madame Bovary. When reading Flaubert's Madame Bovary we do not believe that the story is a true report of known facts. Rather, we imagine that Emma Bovary tries to escape the banalities of provincial life by having adulterous affairs and living beyond her means.
A similar idea has been endorsed in different ways by contemporary upholders of fictionalism about scientific models. Godfrey-Smith submits that 'modelers often take themselves to be describing imaginary biological populations, imaginary neural networks, or imaginary economies' that 'might be treated as similar to […] the imagined objects of literary fiction' (2006, 735). Frigg (2010) developed an original account in terms of Walton's (1990) pretence theory of fiction. Toon (2012) proposed an account of modelling in terms of Walton's (1990) notion of de re make-believe. And Levy (2015) recently advanced a distinct account in terms of Walton's (1993) notion of prop-oriented make-believe. These authors divide between two main distinct approaches. Godfrey-Smith (2006, 2009) and Frigg (2010) uphold an indirect fictionalist account by claiming that model descriptions specify model systems that are akin to fictional characters. Toon (2012) and Levy (2015) advance a direct fictionalist account by claiming that there are no model systems and that models are imaginative descriptions of particular physical systems. As I will argue, neither approach provides a plausible account of the aboutness condition and the epistemic condition.
Walton's aesthetic notion can be developed into a theory of models as representations. Just like a fictional story prescribes imaginings about fictional objects model descriptions prescribe imaginings about hypothetical systems. Model systems do not exist. We merely imagine that they do. Thus, modelling is indirect in the sense that scientists explore and develop model systems in the imag-

ination before asking whether they bear any interesting relation to specific physical systems. The relevant sort of relation (e.g. similarity) must be an intensional relation, i.e. one that is detached from the extension of the terms involved in the relevant claims (Author, forthcoming).
From this fictionalist account of models follows a natural account of the epistemic condition and the aboutness condition. The key to the aboutness condition is in the notion of denotation. Traditionally, denotation is characterised as a two-place relation between a symbol and an object. For example, the proper name 'Napoleon' denotes Napoleon. Likewise models denote their targets. What establishes denotation is a much-discussed topic in philosophy of language. Here I will critically assess two main alternative accounts (descriptivism and direct reference theory) and a plausible contextualist alternative in terms of a user's specific purposes. The key to an account of the epistemic condition is in the notion of propositional imagination, which is an ability to recognize and respond to non-actual scenarios, make assumptions, use symbols and representations of things. In characterising this notion I will rely on previous work I did on the scientific imagination (Author&Author, forthcoming).

Jamie Shaw. The Search for Kuhn-loss: A New Strategy for HPS The notion of 'Kuhn loss' has received extremely little attention in the secondary literature (with the notable exceptions of Midwinter and Janssen 2012 and Chang 2012). The term comes from Heinz Post who used it to describe Kuhn's (1962, 1977) description of how there is a loss of some puzzle solving ability in the initial conception of a new paradigm. This is quite surprising given that it is one of the key points made against the claim that the history of science is linearly cumulative. My paper aims to accomplish three things: I will (i) articulate a clear conception of Kuhn-loss with the help of some historical examples, (ii) demonstrate the theoretical and practice importance of Kuhn-loss (iii) show some of the advantages that the search for Kuhn-loss possesses over other strategies in HPS.
There are several conceptual difficulties with Kuhn-loss. First, it is unclear what Kuhn's thoughts on this matter were since they are not formed in any cohesive manner. For instance, it is unclear whether this means the successive paradigm is initially unable to solve these puzzles or whether it is ever able to solve these puzzles. Kuhn's example of phlogiston theory's ability to account for the qualities of chemical substances suggests that Kuhn-loss at least can be recovered by succeeding paradigms but not that it must be recovered. This gives us a distinction between genuine Kuhn-loss (i.e., loss that has never been recovered) and regained Kuhn-loss. Additionally, while it is clear that since Kuhn himself thought of puzzle solving as one of the primary virtues of a paradigm (and a marker of progress (cf. Laudan (1978)), it is unclear that we should conceive of Kuhn-loss in these terms rather than in the terms of other epistemic virtues (e.g., explanation, prediction, etc.). Some puzzles may not be worth solving from a contemporary standpoint (i.e. what the balance of the four humors is in an individual patient). These kinds of Kuhn-losses are therefore irrelevant from a contemporary perspective. Because of this, I argue that we should partially reformulate Kuhn-loss with other epistemic virtues to give us a notion of Kuhn-loss that is worth regaining.
Several historical examples help develop this concept of Kuhn-loss and its contemporary importance. For example, consider the revival of the cosmological constant (i.e., value of the energy density of the vacuum of space). Einstein's (1917) introduction of this notion was meant to relieve the tension between his view of gravity and the static universe model, but became abandoned due to its conflict with other known phenomena (e.g., red shift). However, the cosmological constant has since been revived in the study of dark matter in an entirely different theoretical context (i.e., in an expanding model of the universe). This example shows how we can conceive of Kuhn-loss in terms of individual principles rather than entire theories. I also consider Copernicus' revival the 'central fire' from Pythagorean thought in his heliocentric model. This shows how metaphysical theses of past scientific theories can be similarly revived (Feyerabend 1975). Or, consider the abandonment of Priestley's version of electrochemistry which was theoretically reconstituted in Lavoisier's combustion theory to regain its explanatory power (Chang 2012). These examples demonstrate how regaining loss has been historically fruitful.
As the aforementioned examples demonstrate, the recovery of Kuhn-loss has aided in the development of science in the past and thus we have strong inductive reasons to continue this activity. For (ii), Kuhn-loss not only has important implications for how we conceive of the history of science, (i.e. as linearly progressive, non-linearly progressive, or not entirely progressive) but also suggests a new avenue for engaging with science. This provides a new set of tasks for historians and philosophers: to find instances of genuine Kuhn-loss, recover them, and apply them to contemporary frameworks. This requires a philosophical engagement with history, which both reveals important philosophical lessons (e.g., incommensurability, the context of discovery/justification debate) and aids progress by proliferating means for improving scientific theories. I go on to make sense of this task and motivate its importance within Feyerabend's account of pluralism (1970, 1975, and 1978).
For (iii), I argue that the search for genuine Kuhn-loss provides a more advantageous way of conceiving of the relationship between the history of science and the philosophy of science. Rather than simply using historical examples to confirm or disconfirm philosophical theories (Laudan, 1981, Worrall, 1989, Psillos, 1999) or using historical examples to illustrate or clarify philosophical theories (Heidelberg and Stadler (2001) and DeWitt (2011)) the search for Kuhn-loss provides a method for engaging directly with scientific practices and aiding in the development of theories. On top of this, the search for Kuhn-loss aids in developing a more comprehensive picture of the development of the history of science by focusing on the conceptual apparatuses of discredited or overlooked theories (cf. Shapin and Schaeffer (1985)) and showing their continuity/discontinuing with successive theories. This demonstrates how the search for Kuhn-loss can directly benefit philosophy, history and science whereas the aforementioned approaches can only directly benefit philosophy of science.

Adam Toon. Extended concepts What are scientific concepts? In this paper, I will offer a new perspective on this question by drawing on recent work in cognitive science and philosophy of mind. Today, scientific concepts are normally taken to be mental representations found in the scientists' head. I will argue that this view is mistaken. In fact, many scientific concepts are what I will call extended concepts: they are realised by interaction between brain, body and world.
Debates concerning the nature of concepts lie at the heart of a range of important epistemological issues concerning the sciences. Perhaps most notable are debates concerning conceptual change and incommensurability. In recent years, a number of authors have shed light on these debates by drawing on research into concepts in cognitive science (e.g. Anderson et al. 2006; Nersessian 2008, Thagard 2012). Traditionally within cognitive science, however, concepts have been held to be mental representations realized by the brain. I will argue that this view is challenged by recent developments in cognitive science and philosophy of mind. In particular, I will suggest that we may gain a new perspective on the nature of scientific concepts by looking to recent debate concerning the extended mind thesis (Clark and Chalmers 1998).
We tend to think of the mind as inside the head. By contrast, the extended mind thesis claims that mental states sometimes extend into the world. Proponents of the extended mind thesis typically focus on (putative) cases of extended belief. In Clark and Chalmers' well-known thought experiment, Otto is an Alzheimer's patient who uses a notebook to compensate for memory loss. Clark and Chalmers claim that Otto's notebook plays a similar role to normal, biological memory. As a result, they argue, its entries count as part of the material basis for Otto's standing beliefs. Let us suppose that we accept the notion of extended belief. Can external devices also extend our conceptual resources? Some work in the literature points in this direction (e.g. Clark 2008). In this paper, I will offer a more direct argument for the existence of extended concepts. To do so, I will draw on a classic study of birdwatching by the sociologists Michael Lynch and John Law (1998).
As Lynch and Law point out, birdwatchers typically venture out into the field with various items of equipment, such as binoculars, spotting scopes, lists and field guides. Most important for our purposes are field guides. The key feature of field guides, of course, is that they are not intended simply to be read in the comfort of an armchair; instead, their illustrations, text and layout are all carefully designed to allow birdwatchers to identify birds out in the field. Drawing on Lynch and Law's account of the practice of birdwatching, I will argue that birdwatchers' use of field guides provides a good example of extended concepts. Just as Otto's notebook forms part of the material basis for his standing beliefs, so the birdwatcher's field guide forms part of the material basis for her conceptual resources. A birdwatcher's concepts are not found (entirely) in her head.
After arguing for the existence of extended concepts, I will show how this idea differs from the familiar content externalism of Putnam (1975) and Burge (1979). I will also consider a number of important objections that might be raised against the notion of extended concepts. First, a critic might argue that concepts cannot be extended since our conceptual resources involve a particular form of representation that is not found in external devices, such as the Language of Thought (Fodor 1975). Second, it might be argued that our concepts display a particular structure, such as that proposed by prototype theorists (e.g. Rosch and Mervis 1975), that is absent in external props like field guides. Third, an opponent might worry that the notion of extended concepts leads to an unacceptable expansion in our conceptual resources. Surely someone doesn't suddenly possess all the concepts of general relativity simply because they borrow a textbook from the library?
I will show that the defender of extended concepts may respond to each of these objections. In fact, I will suggest, rather than leading us to reject the notion of extended concepts, these objections point to the important implications of this notion for our understanding of concept representation, conceptual structure and the limits of our conceptual resources. Recently, a number of authors have argued that the extended mind thesis has important implications for epistemology (Clark et al. 2012; Pritchard 2010; Vaesen 2011). If the argument in this paper is along the right lines, then the implications for the nature of scientific concepts are equally far-reaching.

Dana Tulodziecki. Against Selective Realism(s) The goal of this paper is to test a number of recent selective realist proposals against a case study from the history of medicine: the so-called zymotic theory of disease. I will argue that, just like their ancestors (i.e. proposals by Kitcher, Worrall, and Psillos), and despite the fact that they were designed at least partially with historical cases in mind, these recent proposals cannot account for the change from the zymotic theory of disease to the germ theory. Further, I will argue that these accounts fail not just for the zymotic theory (though that would already be an important result), but that they contain features that call into question their adequacy as workable candidates for selective realism on a more general level.
According to one of the main anti-realist arguments, the pessimistic meta-induction, we have reason to believe that our current theories are just as false as their predecessors. Proponents of this argument draw attention to a list of theories that were once regarded as highly successful, yet ended up being discarded and replaced by radically different ones. Traditional selective scientific realists, such as Kitcher, Worrall, and Psillos, have argued, first, that the anti-realists' list is too permissive, and ought to be restricted only to theories that enjoyed `genuine' success, which, according to realists, consists in a theory's ability to make (use-) novel predictions, i.e. predictions that played no role in the generation of the original theory. Second, in dealing with the remainder of the so diminished list, these selective realists have emphasised the carrying over of stable and continuous elements from earlier to later theories, which are then used to argue for the approximate truth of those earlier theories. Here, the idea is that only those elements essential for the theory's success are carried over, and that, as a result, only they constitute appropriate candidates for realism.
However, these traditional accounts have been criticised on the grounds that they rely on vague notions of essentialness and that, at least partially for this reason, they are still vulnerable to arguments involving cases from the history of science (cf., for example, Lyons 2006 or Vickers 2013). In order to overcome these difficulties, modern selective realists have sought to propose more precise notions of essentialness that avoid both the vagueness of the original proposals and that are also such that the resulting forms of selective realism are immune to the kinds of cases proposed against traditional accounts. The two best worked out proposals for essentialness are those of Vickers (2013) and Peters (2014). According to Vickers, we ought to distinguish between derivation-external and derivation-internal posits, and, further, regard as essential only a specific subclass of the latter. Peters defends a unification account of essentialness and argues "that the essential posits of a theory are those that unify the accurate empirical claims of that theory" (377).
In this paper, I will discuss these two proposals, and argue that neither of them can account for the change from the zymotic theory of disease to the germ theory. The zymotic theory was one of the most sophisticated and popular versions of the mid-nineteenth miasma theory. According to the zymotic theory, diseases occur as a result of introducing into the body various zymotic materials, either through direct inoculation or through inhalation after being dispensed in the air. Essentially, these zymotic materials were thought to be putrefying organic matter that would communicate its process of decomposition to pre-

existing materials in the victims' blood where it would act in a manner similar to ferment, thus causing diseases. Due to its analogy with fermentation, the zymotic theory was able to draw on some of the most successful science at the time, such as Liebig's chemical theories, thereby allowing it to propose highly detailed mechanisms about the exact manner of disease causation.

I will show that the zymotic theory was highly successful and made a number of use-novel predictions, some of them of striking quantitative precision. Thus, it was successful in the realist's sense; however, despite its successes, it turned out that its central theoretical elements – zymes and miasmas – turned out not to exist. Neither are there other candidates for continuity between it and the germ theory. However, as I will argue, this case is not just a problem for the traditional selective realists, but also for modern selective realists such as Vickers and Peters, since zymes and miasmas ought to be regarded as essential on both their accounts. More specifically, zymes and miasmas ought to count as essential derivation-internal posits for Vickers, and, further, also meet Peters' three unification criteria. Thus, modern selective realist accounts – even those designed specifically with cases supporting the pessimistic meta-induction in mind – fail to meet challenges from the history of science. Although the zymotic theory constitutes only one case, I will argue that this example supports a broader conclusion, since it highlights more general features of the accounts of Vickers and Peters that suggest that this case will not remain the only one that is troublesome for them.

John Wigglesworth. Logics as Scientific Theories of Consequence Logic is the study of consequence, the study of what follows from what. Different logics give us different theories about what follows from what. For example, classical logic says that P follows from ~~P, while intuitionistic logic denies this. These two theories disagree about what follows from what, and so they give different, competing theories of consequence. If logics are theories of consequence, it is natural to think that they are scientific theories of one sort or another. And the study of the nature of scientific theories falls to the philosophy of science. This paper explores the treatment of different logics as giving competing theories of consequence from the perspective of the philosophy of science.

The literature on scientific theories is split roughly into two camps: syntactic accounts and semantic accounts. Syntactic accounts take theories to be sets of sentences in a particular language. Semantic accounts take theories to be sets of models, without any syntactic or linguistic component. Problems arise for both views.

Many of these problems are concerned with the identity conditions for theories. In both cases, theories are set-theoretic structures — either sets of sentences or sets of models. Identity conditions for sets are given by their extensional nature: two sets are identical if and only if they contain exactly the same members. One can show, however, that the same set of sentences may give rise to two intuitively distinct theories. This result raises a challenge to the syntactic view, which would be forced to judge these intuitively distinct theories as identical. The converse problem arises for the semantic view: there are theories that are intuitively the same, but which can be identified with different classes of models (Halvorson 2012).

When taking different logics to be competing scientific theories of consequence, one can raise the same kinds of counterexamples mentioned above to both the syntactic and semantic approaches. We argue that a semantic approach to scientific theories is preferred in the case of philosophical logic, and that the standard counterexamples to the semantic view can be resolved in this case.

The general project of taking logics to be scientific theories of consequence has several further objections to answer to. We consider one such objection, which is based on the idea that standard approaches to scientific theories are ill suited to the task of articulating different logics as theories of consequence. The problem is that scientific theories should contain some logical apparatus from the beginning, and so they already assume some notion of consequence. For example, on the syntactic view, a theory is not just any arbitrary set of sentences, but a set of sentences together with the consequences of those sentences. Theories are closed under the consequence relation. Any theory of the notion of consequence will therefore assume some particular consequence relation already. And so the view that logics are competing theories of consequence will involve an inherent circularity.

We show how to defuse this objection by arguing that the logic of a scientific theory is simply another part of the theory (this is easily done on the semantic view of theories). The logical principles of a theory do not have any special status, over and above the theory's more "empirical" assumptions or consequences. A theory's logical principles are simply another part of the theory, just as open to reflection and revision as any other part. Revision of logical principles is accomplished through familiar scientific methodology, which evaluates theories according to certain criteria, or in terms of how they exhibit certain theoretical virtues. These include simplicity, strength, elegance, explanatory power, and coherence with what we already know. The last of these is particularly important, as it measures how one's theory fits the available data. To conclude, we consider what phenomena comprise the relevant data when the scientific theory under consideration is a theory of logical consequence.

J.E. Wolff. Why did additivity cease to be the central element in the foundations of measurement? In the key writings on the foundations of measurement from the late 19th to the early 20th century (e.g. [1],[2],[3]), additivity was the central criterion for establishing appropriate axioms for measurement. While authors disagreed about which axioms of additivity were relevant, and how best to implement these axioms in the foundations of measurement, there was little doubt that additivity was a necessary condition for a property to count as measurable, or at least for the property to count as an extensive quantity. This changed fairly abruptly around the middle of 20th century, and today the most comprehensive treatments of the foundations of measurement have relegated additivity to secondary status. A key figure in bringing about this shift in attitude was S.S. Stevens, whose article on the "scales of measurement" [4] became hugely influential in shaping what we now think of as the foundations of measurement.

My topic in this paper is why and how this change came about. In particular, in what sense, if any, can this change in the foundations of measurement be regarded as a Kuhnian paradigm shift?

Stevens' own dismissal of additivity provides a good starting point: "Too much measuring goes on where resort can never be had to the process of laying things end-to-end or of piling them up in a heap." (Stevens, 1946, 680). An account of the foundations of measurement based on axioms of additivity turns out to be too restrictive to capture everything one might want to call measurement in science, and should hence be abolished. Stevens' alternative suggestion is that we should adopt a wide notion of measurement, and put restrictions only on the statistical inferences we may draw based on the data arrived at through different measurements. Stevens' position appears as that of a pragmatic social scientist, who has to defend the methodology of his discipline against (dogmatic) physicists. For Stevens' paper is not merely a methodological reflection for its own sake, it is a response to disagreements among physicists and psychologists in a committee set up to debate whether sensation was measureable. What we have here, then, is a direct confrontation between different scientific research fields over fundamental issues of methodology.

One prominent member of the committee on the side of the physicists was Norman Campbell, who had written a fairly comprehensive treatment of measurement in physics in the early 1920s based on axioms of additivity [3]. A closer look at both Campbell and Stevens points to a second reason, not purely pragmatic this time, for why additivity lost out. Even Campbell himself had already recognized that some quantities fail to meet the requirements of additivity and only allow for ordering. His response had been to relegate those quantities to derivative status. Stevens' strategy instead employs different mathematical tools, notably group theory, to display the differences in status. The second reason for the disappearance of additivity might hence be the availability and use of more modern mathematical tools, which had either not been available during previous work on measurement, or which in any case had themselves seemed insufficiently secure.

So far then we have pragmatic reasons and the availability of alternative formal tools to account for the change in the importance of additivity. But while the tools employed by Stevens provided a reason to think that a looser notion of measurement might be feasible, it wasn't until the discovery of conjoint measurement that additivity could be dropped as a necessary condition even for extensive quantities. The development of conjoint measurement demonstrated that three suitably related attributes could be shown to have continuous quantitative structure even in the absence of any concatenation operation that could be interpreted as addition [5]. The possibility of conjoint measurement seemed to show conclusively that additivity was not a necessary condition for the measurement of extensive quantities.

I conclude that additivity fell out of favor as a necessary condition for measurement on the one hand because new scientific disciplines demanded a relaxation of apparently overly restrictive standards applicable to physical quantities, and on the other hand because new formal techniques were developed that allowed for a characterization of extensive quantities not based on the axioms of additivity.

This episode in the history of measurement theory is philosophically interesting in many respects. It shows how a conceptual change—in our understanding of the concept of extensive quantity—is accomplished not through new data or empirical results, but arises from new demands in emerging scientific fields, and the development of new formal methods. In this way the switch from axioms of additivity as a basis for foundations of measurement meets Kuhnian conditions for paradigm shifts: a new paradigm, in this case new formal tools, has to be available before the shift can be completed. A result of this change, a sociological shift in science occurs as well: the foundations of measurement emerge as their own field of research, after having been regarded firmly in the domain of physics. In another respect the shift is remarkably un-Kuhnian, however: since measurement concerns not just a single discipline, it would be misleading to characterize the shift in understanding of measurement and quantities as a paradigm shift within a single field of research. Instead we see the emergence of an independent area of research out of the conflicting demands of different scientific disciplines.

K. Brad Wray. What Support Do the Theoretical Values Offer the Scientific Realist? Many realists argue that the theoretical values are reliable indicators that a theory is likely approximately true with respect to what it says about unobservable entities and processes. Anti-realists disagree, claiming that these values are not reliable indicators of theoretical truth. I argue that these values are not capable of supporting the sorts of claims that realists lead us to believe they can support. Scientists cannot infer that theories that embody these values are likely true or approximately true. And there is no reason to believe that these values systematically track theoretical truth. I also argue that the theoretical values fail to provide scientists with the practical guidance realists suggest they can offer in choosing which theory to work with.

First I argue that realists cannot get the evidential leverage they need from the theoretical values, even granting that the theoretical values are correlated with theoretical truth. When a scientist makes an evaluation of the simplicity or scope of a theory, such a judgment is not a categorical judgment, like the judgment that a proposition is true. Such a judgment is comparative. The scientist judges that one theory is simpler than another theory. On the basis of this sort of judgment, the scientist is not in a position to infer that the simpler theory is true, or even approximately true. All that such judgments yield is an ordinal ranking of the competing theories. When one theory is judged to be simpler than another, there is no fixed benchmark of simplicity against which this judgment can be measured. So such evaluations, insofar as they support judgments about the truth or approximate truth of a theory, merely support the claim that the one theory is closer to the truth than the other. But that is quite a different matter than inferring that the one theory is likely true, or even approximately true. Indeed, from an evaluation of two competing theories with respect to their relative simplicity, we are in no position to know how far either theory is from the truth. All we can legitimately infer is that the one theory is closer to the truth than the other. But both theories could be significantly far from the truth.

I then consider the extent to which the theoretical values can guide scientists in their pursuit of better theories. Realists assume the rational action is to work with the theory that is closest to the truth, the simpler theory, for example. But even this is inference is fallacious. Even after it is determined that one theory is simpler than another theory, thus closer to the truth, one cannot infer that accepting and working with the simpler theory is more apt to lead scientists to the truth or even closer to the truth in the long run than working with the more complex theory. Features of the simpler theory may prove to be impediments to further improvements. Once scientists recognize that they are working with imperfect theories, they should realize that working with a theory that is closer to the truth is not necessarily the best path to follow in order to get a better theory. Working with the more complex theory may be the more expedient path to a better theory.

Consider the operation of natural selection in the biological world. In the biological world, a species can get caught in a local maximum, a position that is inferior to the global maximum, but superior to any position around it. Once a species has reached a local maximum, any move toward the global maximum would require a step away from the local maximum, in the direction away from fitness, at least temporarily. But natural selection is shortsighted. The processes at work in natural selection will not subject a species to short term loses in fitness for long term gains in fitness.

Similarly, consider a case where scientists are choosing between two competing theories, T1 and T2, where T2 is deemed to be closer to the truth as determined by the theoretical values. Choosing to work with a particular scientific theory, T2 rather than T1, may prevent scientists from getting to an even better theory, T3. Features of the one theory deemed to be superior (T2) may prove to be serious impediments to future improvements in a way that comparable features in its weaker competitor (T1) are not. But given that the realist assumes a strong connection between the theoretical values and theoretical truth, the realists' advice to scientists will be to work with the theory that embodies the theoretical values to the higher degree. This strategy, I argue, will make scientists vulnerable to getting caught in a local maximum. Sometimes it will be in the interest of scientists to work with a theory that is further from the truth (T1) than some existing competitor (T2). Doing so may lead to an even better theory (T3) in the long run.

This argument is not a new version of the Underdetermination of Theory Choice Evidence, nor is it based on an extreme form of skepticism about induction. Rather, it is based on the realization that the path to ever-better theories is not necessarily linear or progressive.

My arguments suggest that realists misunderstand the role that theoretical values should play in science. The theoretical values fail to both (i) provide warrant for the inferences that the realists want to draw, and (ii) solve the practical problem of determining which theory a scientist ought to work with. When scientists appeal to the theoretical values to evaluate competing theories, they merely get an ordinal ranking, not the sort of ranking that can support the inference to the likely truth of the superior theory. Further, if scientists are aiming to develop better theories, given a set of theories to choose from, working with the theory that embodies the theoretical values to the greatest extent may not be the most expedient way to do this.

Nicolas Wüthrich. The puzzle of the persistence of non-robust toy models Simple, highly idealized models can be found across the sciences: The Lotka-Voltera model of predation in biology (Lotka 1956), Shelling's checkerboard model of segregation in sociology (Schelling 1978), or the Arrow-Debreu general equilibrium model in economics (Arrow and Debreu 1954) are just some prominent examples. Recently, there has been growing interest in the epistemic function which these so-called toy models play (see for example Batterman and Rice 2014; Grüne-Yanoff 2009; Weisberg 2006).

In this paper, I identify and dissolve a puzzle regarding a subset of this model class: There are toy models which are persistent although their key results are non-robust, they do not display heuristic value for generating new hypotheses or models via de-idealization, and they have a relatively poor predictive track record. To do so, I argue that the consensus in the literature on toy models, namely that non-robust toy models have to be revised substantially or to be rejected altogether (see for example Hands 2016 and Cartwright 2009), needs re-thinking. I defend the view that toy models can have epistemic value despite the fact that their key model results are not robust to changes in the auxiliary assumptions.

I start by introducing in some detail a case study from macroeconomics. Dornbusch's famous overshooting model for exchange rates (Dornbusch 1976) aims at explaining why exchange rates between countries are more volatile than underlying fundamentals such as price, output, or money supply levels suggest. Dornbusch showed that exchange rate overshooting can be present without assuming irrational behaviour of agents but instead by presupposing that the goods and asset market in an economy do not adjust equally fast to an exogenous shock. I show that this model is an instance of the puzzle: The model persists, particularly in policy-making circles, despite the facts that its key model result is not robust to changes in auxiliary assumptions, its empirical track record is widely criticized, and it does not have heuristic value as state of the art models of international monetary theory cannot be arrived at via de-idealizing the Dornbusch model.

In a next step, I argue that a closer look at the economic practice reveals a starting point for identifying potential epistemic functions of these non-robust toy models. The Nobel laureate Paul Krugman claims that models are "crucial aid to intuition" (Krugman 1998, 1834). This practice, which is an often echoed view of toy models as contributing to the sharpening or exploring of our intuitions before one turns to more complex although empirically more adequate models, needs to be made more precise.

I start by showing that two salient strategies to make this talk more precise fail in light of the non-robustness of model results. First, toy models which are instances of the puzzle do not sharpen our intuitions by giving us information about constraints on relations between quantities in the target system. The lack of robustness is a good reason to suspect that the model equations and derived model results do not represent constraints in the target system. Secondly, toy models do not sharpen our intuitions by changing our credences in impossibility claims (e.g. exchange rate overshooting cannot occur without some irrationality of agents). This account of learning with minimal models has attracted quite some attention in the recent literature (Grüne-Yanoff 2009, 2013; Fumagalli 2014). Grüne-Yanoff (2009)'s account fails to articulate conditions under which the possibilities, which are displayed by the model results, can be viewed as relevant possibilities for a target system. Furthermore, the lack of robustness of model results is an indication that the model results are not relevant possibilities in relation to the target system.

I, thirdly, show that there is an alternative way of explicating the "sharpening of intuition"-strategy. I argue that toy models sharpen our intuitions, and, hence, provide epistemic value, via two heuristic functions. Let me be clear that I deliberately attach epistemic value to heuristic functions of models. Non-robust toy models can have two heuristic functions.

To start, non-robust toy models can identify a stylized problem with agenda-setting character for a discipline. In the case of Dornbusch's overshooting model, this stylized problem consists in determining the effect of a completely unexpected but permanent rise in the money supply. Studying this problem can lead, at a further stage of the development of the discipline, to more accurate explanatory and predictive models, and, hence, epistemic insight.

Furthermore, non-robust toy models can provide explanatory templates for viewing causal processes. The explanatory template provided by the Dornbusch model is characterised as 'non-smooth behaviour with rational expectations'. This explanatory template does not allow us to debunk the less abstract impossibility claim 'exchange rate overshooting cannot occur in a rational expectation framework' (as Grüne-Yanoff would suggest). Instead, it makes the very concept of non-smooth behaviour in a rational expectations framework available for the discipline. The case study shows that this explanatory template was important for developing empirically more accurate state of the art models in international macroeconomics. Hence, the epistemic value of a toy model, which is given by providing an explanatory template, has to be viewed in relation to the success of models which are developed later in a field. Importantly, the state of the art models in exchange rate economics (the so called New Open Economy Macroeconomics) cannot be viewed as de-idealizations of the Dornbusch model given one opts for the widely shared account of de-idealization proposed by McMullin (1985).

In summary, I suggest to not assess toy models directly in relation to target systems of interest but indirectly with respect to the adequacy of models which are developed later in a field and which are linked via an explanatory template or a stylized problem to the toy model. This shift of perspective allows not only shedding light on epistemic functions of non-robust toy models but also defining new criteria of adequacy for this type of modelling practice.

Robin Zheng. Responsibility, Causality, and Social Inequality Apparently empirical disputes about the causal explanations of social inequalities—in particular, whether they are attributable to individual dispositions or background structural factors—actually rest on conflicting normative expectations concerning the distribution of powers and roles that ought to have prevented some event or state. These normative expectations, in turn, depend on cultural and moral values derived from individuals' backgrounds and lived experience. Thus, I argue, philosophers who work to reshape such normative expectations and values also work to restructure what count as acceptable causal explanations of, and hence interventions on, existing social inequalities. In other words, philosophers play an indispensable role in diagnosing and addressing the deep disagreements about causal explanation that often stymie public discourse and policy around social inequalities.

I begin by drawing on a familiar insight from feminist philosophy of science: because distinctively philosophical assumptions undergird all scientific inquiry. I illustrate this with an extended example of the intertwined philosophical and social psychological research surrounding the problem of poverty in the United States. Research on causal attributions for poverty since the 1970s has demonstrated that people give causal explanations of poverty that fall into (at least) two distinct categories: individualistic, which locate the causes of poverty in the dispositions of poor people, and structural, which locate them in broader social and institutional factors. This distinction between individualistic and structural attributions is especially significant because it reflects one of the core notions of social psychology: the "fundamental attribution error," or the psychological tendency to view a person's behavior as caused by stable dispositional traits of the person herself, rather than by contingent features of the situation in which that person is acting. One of the most notable things about the fundamental attribution error is that research has uncovered significant variations across social groups: women, lower-income, non-white, low-education, and liberal subjects are more likely to endorse structural (situational) explanations while men, middle-income, white, moderately educated, and conservative subjects are more likely endorse individual (dispositional) explanations. These findings provide striking support for the claims of feminist epistemologists that structures and processes of knowledge-making are shaped by social location.

In the second half of the paper, I develop the claim that disagreements about causal explanation are not factual, but moral disagreements. Here, I draw on the work of Joel Feinberg and Marion Smiley to show that causal explanations about an event or state depend on beliefs about how it could have been prevented, which in turn depend on beliefs about social roles and expectations about who or what in the community possesses or ought to possess such powers. Recognizing these implicit assumptions about social roles expectation, which in turn are structured by power, allows us to see how questions about causality depend irreducibly on moral and political questions, and how the social science of poverty and inequality depend on assumptions grounded in (moral) philosophy. Returning to the problem of poverty, I apply this framework toward understanding Iris Marion Young's analysis of the rhetoric of "personal responsibility" surrounding the problem of poverty.

I conclude by discussing two different ways in which philosophers can contribute, on the back end, to science and to real-world problem-solving. On the one hand, philosophers can involve themselves in deeper critical engagement with empirical research. For example, debates about construct validity—that is, whether a given measure really does measure what it is supposed to, or the put it another way, how to operationalize some concept like "well-being" or "poverty" in terms of measurable empirical indicators (cf. Thomas Pogge's work)—can be usefully informed by the conceptual analysis that forms the bread-and-butter of analytic philosophy. On the other hand, philosophers can also work to change the normative expectations that underlie our causal explanations of social problems. We can also view the work of philosophers like Peter Singer and Thomas Pogge—who have argued forcefully for expanding our ethical obligations to include duties to the global poor—as doing just this. Certain causal explanations, such as, say, "Wealthy nations cause global poverty", become intelligible once we hold them to the moral expectations advocated by these and other moral philosophers.