INVITED SPEAKERS

Which mechanicism for mechanisms in medicine?

Raffaella Campaner

Back in the early Eighties, mechanicism as the way to grasp causation and scientific explanation was advocated by Wesley Salmon as part and parcel of a revival of causation in a probabilistic framework. Ever since then, the notion of mechanism has enjoyed great success in contemporary philosophy of science, so much so as to be at the core of a proliferation of philosophical theories and – more or less critical – positions on possible applications in a wide range of fields. The paper aims to discuss the development of the mechanistic approach from the mechanistic picture of "the causal structure of the world" (Salmon, 1984) up to current different versions of the "new mechanical philosophy" (Glennan, forthcoming). The paper evaluates the recent shifts the notion of mechanism has had, referring, more specifically, to causal discourse in psychiatry and discussing what mechanisms can and cannot do in that context in the light of stances from the philosophy of science and from within psychiatry itself.

Daniel Commenges

Causality: marginal effects, conditional effects and mechanisms

I first present some general remarks about causal questions. The counterfactual approach and an approach based on physical laws are presented. A mathematical theory of dynamical influences can be developed and such dynamic influences can be represented graphically. Within this theory it is shown that randomized trials can estimate the marginal causal effect of a treatment. However, there are several limitations to randomized experiments. We need to know mechanisms, in the first place to select therapeutical strategies that will be tested by randomized trials, but also to develop adaptive strategies. So, we need to develop mechanistic models of pathological conditions. This can be done at different levels (populations of subject, subjects, populations of cells, cells), with different degrees of precision, and the models can be qualitative or quantitative. These questions can be illustrated using HIV infection and AIDS. We can consider infection at the level of the population, of the subject, of populations of cells, of cells. Quantitative models of interaction between population of viruses and populations of cells of the immune system have been developed. These could be used for designing adaptive strategies for tuning the dose of antiretroviral drugs. Also treatments using interleukins have been proposed to restore the immune system in HIV infected subjects. Here also, models have been developed to adapt the protocol of administration of interleukins. Adaptive strategies have always existed but they can become more precise thanks to mathematical/statistical models.

The Precautionary Principle meets the Hill Criteria of Causation: A Case Study of Tuberculosis among Gold Miners in South Africa

Daniel Steel and Jessica Yu

This article examines the relationship between the precautionary principle and the well-known Hill criteria of causation. Some have charged that the Hill criteria are anti-precautionary because they are strongly inclined towards false negatives in multi-causal contexts typical of environmental and public health issues. However, we argue that without guidance on how to interpret and weight the criteria, no such claims can be supported. We illustrate this argument with a case study of tuberculosis among South African goldmine workers wherein different weightings of the Hill criteria suggest different conclusions about mechanisms, which in turn has implications for an ongoing class action lawsuit brought by former miners.

Does mechanism research or clinical observations generate more medical discoveries? Analysis of the meta-research data

Jeremy Howick

Broadly speaking, there are two overlapping methods for discovering new treatments: basic science research—usually into underlying mechanisms) and clinical observations. Philosophical proponents of mechanisms support have claimed that mechanism research is a superior strategy for generating medical treatment discoveries. In contrast, medical researchers have recently claimed that much mechanism research is 'wasted' because it leads to new discoveries far less often than mechanism proponents claim. In this paper I examine the empirical evidence for the relative success of basic mechanism research compared with clinical observations for generating treatment discoveries. I find that the empirical evidence is inconclusive and highly heterogeneous. As such, it is currently impossible to say which strategy, if any, is more successful.

Evidence-Based Medicine and 'Patho-Physiologic Rationale': It's the evidence, stupid!

John Worrall

OF COURSE there is stronger evidence for the effectiveness of some treatment T for condition C if, alongside evidence from clinical trials, there is evidence from underlying physiology/biochemistry that T positively affects C. So what were the founders of EBM thinking of in seeming to downgrade 'patho-physiologic rationale' as a source of evidence? And what is the basis of the ongoing dispute between Jon Williamson and Jeremy Howick on the evidential role of 'mechanisms'? This paper investigates and in the process looks in detail at particular cases involving issues of 'patho-physiologic rationale' that influenced the initial EBM position (notably grommets for glue ear and the suppression of ventricular ectopic beats for reducing risk of heart attack); as well as more recent cases where a proper weighing of the strength of clinical trial evidence for the effectiveness of a treatment depends crucially on knowledge of underlying physiological/biochemical facts (notably the case of bisphosphonates as a treatment for osteoporosis). The basic position advocated should be

familiar from my previous work: not so much EBM^{+} as PEBM – Properly Evidence-Based Medicine.

In Defense of Methodological Mechanism: the Case of Apoptosis

Stathis Psillos & Stavros Ioannidis

The aim of this paper is to propose a minimal way of understanding mechanisms and to argue that this conception is sufficient in order to have an illuminating account of scientific practice. We will focus our attention on a central example of a biological mechanism, the mechanism of cell death, known as *apoptosis*. In particular, we are going to use this case in order to advance the philosophical position we call 'methodological mechanism': the claim that to be committed to mechanism is to adopt a certain methodological postulate, i.e. to look for causal pathways for the phenomena of interest. Methodological mechanism and the deflationary account of mechanisms that it incorporates differs from recent accounts of mechanisms in that it is ontologically non-committal; moreover, our argument will unravel some limitations inherent in any attempt to extract metaphysical conclusions from scientific practice.

CONTRIBUTED PAPERS

Kinds of Mechanistic Evidence

Daniel Auker-Howlett

One problem facing the integration of mechanisms into clinical decision mak-ing is how to know when we have a complete mechanistic explanation for a phe-nomenon. Using incomplete mechanistic evidence to make decisions on treat-ments have in the past led to adverse outcomes. A challenge then is to identify when we have a complete mechanism. But rst, to arrive at any mechanistic explanation, we must discover it. Craver (2007) and Craver and Darden (2013) have argued that this discovery process is made easier by the use of a reasoning strategy where evidence applies constraints on a space-of-possible-mechanisms (SOPM). The aim is to arrive at a how-actually mechanistic explanation that correctly describes the mechanism in question. They state that do so evidence from experimentation applies various types of constraints on any possible mech-anism, such as locations of components, or timing of a process. In this way, evidence just is the output of experiments and the justi catory work is per-formed by type of constraints.

Contra this evidential homogeneity, I argue instead that evidence is varied in kind and is picked out by the features of the evidence itself. To do so I make a distinction between evidence derived in a search for mechanisms, or in a mechanistic framework, and evidence which is not. For evidence derived from outside of a mechanistic framework, the evidence is used wholly inferentially as the original empirical data are not used to justify constraints. In doing so I argue for a functionalist conception of evidence where it is not just empirical data that count as evidence. I instead argue that how the evidence is utilised to provide certain explanatory constraints on the SOPM is the condition that di erentiates ve kinds of evidence. For evidence derived from outside of a mechanistic framework, the di erentiation is down to: whether the evidence consists of a phenomenal characterisation providing constraints on potential ex-planations; or, whether it is highly-con rmed theory of physical, chemical or bio-chemical principles that limit the space to possible explanations. For evi-dence derived from within a mechanistic framework the distinctions are three-fold. Observational evidence can be either indirectly-seen to weakly-plausible, or directly-seen as strongly-plausible explanations. Simulation is a nal distinct kind of evidence, where computational or synthetic modelling of a system can provide either weakly-plausible constraints when not supported by OE; or actual constraints when the model is constructed from observational evidence that has produced a strongly-plausible explanation.

My conclusions expand on the characterisation of SOPM to include these kinds of evidence. This expansion means that we can more easily identify when we have a complete mechanistic explanation, as certain kinds of evidence entail greater justi cation. I propose that only when we have a complete network of kinds of evidence can we identify whether we have a complete mechanistic explanation.

Breastfeeding and Child Health: Why you may not want to randomize even if you could.

María Jiménez-Buedo. (with Héctor Cebolla-Boado and Leire Salazar)

During the last couple of decades the list of health benefits associated with breastfeeding (both for mother and child), as widely publicized by health authorities, seemed to be growing yearly. The most cited benefits for full -term children include: lower risk of Sudden Infant Death Syndrome (SIDS); lower incidence of food allergies, asthma and eczema; lower incidence of obesity and type II diabetes; lower incidence of gastrointestinal, respiratory and ear infections (Acute Otitis Media); and at the cognitive level, higher I.Q. and even lower incidence of Attention Deficit Hyperactivity Disorder are often claimed.

Yet, after years of public campaigns promoting universal (and exclusive) breastfeeding, the idea that "breast is best" has made its way into shared public wisdom right at the time when the scientific consensus over the question is beginning to falter. An increasing number of studies are acknowledging the problematic character of much of the accumulated evidence regarding the health benefits of breastfeeding. Much of this evidence is now considered to suffer from *selection bias* problems: the association between breastfeeding and child development or health outcomes would be explained by the characteristics of the families that choose to breastfeed (generally, in developed countries, households with more resources or more information). It would be these characteristics, rather than the milk itself, that would produce the benefits for the child. Because of this, studies that use sibling studies (with data from families in which there are both breastfeed and formula fed children) (Colen and Ramey 2014), or studies that control for parents characteristics through fine grained variables such as parenting styles, rather than merely income (Gibbs and Forste, 2014), have debunked already many of the advantages associated to breastfed children.

Together with the more careful use and collection of observational data, there is also the important contribution of Kramer's PROBIT (the Promotion of Breastfeeding Intervention Trial), a cluster-randomized trial conducted in 1996-1997 in thirty-one maternity hospitals in Belarus, including as many as 17.046 mother-infant pairs followed through time. A subset of these hospitals (16) were subject to an intervention modeled on the Baby-Friendly Hospital Initiative of the WHO (emphasizing assistance with initiating and maintaining breastfeeding and lactation), whereas the rest served as controls, by letting them continue usual infant feeding policies (Kramer et al. 2001). As the first and to date only large scale breastfeeding trial, Kramer's PROBIT has already contributed to the sobering of breastfeeding enthusiasts by showing that breastmilk has no effect on diabetes and obesity prevention, and only very modest results on the prevalence of gastrointestinal tract infection and atopic eczema. Also, it has shown to have had a very modest increase in the IQ (3 points) of breastfed children when measured at the age of 6.5 years (Kramer 2008), and the evidence on whether these

results hold in children over 13 years of age has not yet been published. These overall modest results are nevertheless frequently quoted and interpreted by breastfeeding advocates and public health officials as proof that breastfeeding is beneficial even under the rigorous test of RCTs.

Our paper takes issue with the limits of randomization in the study of breastfeeding and its related health outcomes. Our strategy is two-fold. First, we discuss some of the physiological determinants of breastfeeding success (for both mother and child). We argue that breastfeeding success plausibly involves unobservable characteristics of the newborn that can correlate with the infant´s overall health, and that these in turn are likely to interact with the mother's determination to breastfeed her child (and thus with breastfeeding duration). We show how under these assumptions, random allocation to a breastfeeding promotion program can bias the data in ways that are equivalent, empirically, to the selection bias that randomization is supposed to counter.

Second, we argue that a cost-effective alternative to RCTs consists of the analysis of observational data from social contexts were breastfeeding is not normativised (i.e., where households with more parenting skills do not self select into breastfeeding). To test this claim, we analyse data from the China Family Panel Study (CFPS), a large-scale representative sample of Chinese households produced by the University of Beijing since 2010 (Xie & Hu, 2014) and show how in the Chinese social context, in which breastfeeding practices are not associated univocally with parenting styles, breastfeeding does not appear to impact positively neither children's health nor cognitive outcomes.

Predicting the Results of Medical Interventions: When Mechanistic Models Misfire

Jonathan Fuller

Mechanistic models of medical interventions often misfire, generating poor predictions of treatment effectiveness and harm. For example, many prospective treatments that appear promising on the basis of preclinical research (including our mechanistic understanding of how the treatment works) flop in premarketing clinical trials. Moreover, sometimes drugs that are used off-label based on mechanistic rationale do not produce the hoped-for benefits. There are several dramatic examples of failed mechanistic predictions in medicine that led to harm on a grand scale; many of these examples are used by the evidence-based medicine (EBM) community to justify placing mechanistic reasoning at the bottom of an evidence hierarchy (Howick 2011). Moreover, mechanistic treatment predictions are often imprecise in the face of treatment effect heterogeneity: we predict that the intervention will benefit some patients, but not which patients will benefit.

In this paper, I provide a general account of mechanistic prediction for medical interventions to help understand the above problems. I start by characterizing mechanistic treatment predictions as a process of reasoning through a model of either (i) an intervention mechanism, or (ii) an intervention on a mechanism; I distinguish between these two distinct

kinds of models. Whenever a mechanistic prediction goes wrong, we could in principle blame: the predictor reasoning through the model (for reasoning poorly), the mechanistic model (for being inadequate), or the mechanism (for being complex or unpredictable). On my account, so long as the predictor is reasoning validly we can understand prediction problems as originating in the mechanistic model, due to either model incorrectness, model incompleteness, or model abstractness.

I assimilate several notable examples of failed treatment predictions under my account, including antiarrhythmic drugs and cardiac death (Howick 2011), the Bangladesh integrated nutrition program and infant malnutrition (Cartwright 2012), and low tar cigarettes and lung cancer (Broadbent 2013). My account also helps clarify several problems raised by philosophers of science, including the problem of paradoxical mechanisms (Howick 2011) as well as the Masking Problem and the Complexity Problem (Clarke et al. 2014). Using a basic framework for representing models of medical interventions, I show that not all instances of model incorrectness, model incompleteness and model abstractness impede successful mechanistic prediction.

A general account of mechanistic prediction in medicine could help in the development of schemes for assessing the strength or quality of mechanistic predictions. Crucially, it is not just the quality of our evidence of mechanisms that is relevant, but also the three properties of mechanistic models – correctness, completeness and abstractness - that my account explores.

Varieties of Process Tracing and Methodological Issues

Virginia Ghiara

The aim of this paper is to examine how mechanistic evidence can be collected in medicine, and what methodological issues can arise when different methods of collection are employed. The main focus is on a method for identifying causal mechanisms with a long tradition in the philosophy of causality: process tracing (PT).

Recently, Beach and Pedersen (2013) advanced the debate on PT by identifying three variants of this method: theory-testing PT, theory-building PT, and explaining-outcome PT. Although they considered only PT in the social sciences, it is possible to extend their proposal to medicine as well. Theory-testing PT deduces a mechanistic hypothesis from the existing literature and then it tests whether there is evidence that the hypothesized causal mechanism actually operates as expected. Theory-building PT is used when it is known that a correlation exists between X and Y, but the potential mechanisms linking the two has not been discovered. The aim is to trace a causal mechanism that is expected to be present across a selection of cases. Explaining-outcome PT seeks to trace the intricate conglomerate of case-specific causal mechanisms producing the specific outcome in question.

This paper extends this classification by distinguishing between three methodological approaches that can be used when PT is performed: case-study PT, large-scale PT, and experimental PT. Case-study PT is performed when scientists select a representative case (or

a group of representative cases) of the target population and collect mechanistic observations. An example is a study performed by Abdel-Karim et al. (2007), who looked at autopsies of a group of patients to ascertain the direct cause of death of individual hospice inpatients.

Large-scale PT, rather than being focused on a single representative case, involves a large class of cases and it is based on the analysis of data related to this cohort. The EnviroGenomarkers project is a case in point: it investigates the effects of environmental exposure on various diseases using evidence from six European countries (Russo and Williamson 2012).

Finally, experimental PT requires both the selection of a model population, and the design of a specific experiment through which the mechanism can be traced. A typical case is a piece of research proposed by Fennell and Brown (2001) on ethylene oxide: by using mice as a model population, scientists were able to trace the mechanism linking inhalation of ethylene oxide to the development of different forms of cancer.

Each of these variants entails particular implications for research design. The selection of a *representative* population is essential in order to perform correctly case-study PT, and the role played by idiosyncratic characteristics has to be studied carefully. On the contrary, in large-scale PT, scientists are confident that all the *average features* of the population under study are taken into account as possible causal relevant factors. Finally, experimental PT requires not only the development of an *adequate experiment*, but also the choice of an *adequate model population* sharing the mechanism under study (or its relevant parts) with the target population.

A Role for Mechanistic Reasoning in Medical Decision Making

Ashley Kennedy

A recent and prominent theme of discussion in the philosophy of medicine literature concerns the evidential role that mechanistic reasoning can (or should) play in medical decision-making. Some (such as Howick 2011) have argued that this role should be considered secondary to research evidence derived from randomized controlled trials (RCTs). Others have argued against this claim (Clarke et al. 2014). Here we suggest an important role of mechanistic reasoning that has been left out of this discussion entirely, namely, its importance in the patient's decision-making process – in deciding whether or not to take certain drugs.

Case: chemical treatment for schizophrenia

The signs and symptoms of schizophrenia are often debilitating and include, but are not limited to, delusions, hallucinations, and disorganized speech. Treatment of the condition is multi-faceted but nearly always includes a chemical component, most often an antipsychotic medication, such as questiapine, risperidone, or aripiprazole. Clearly, as with all cases of treatment decision-making, the risk of side-effects from these medications must be weighed

against the potential benefit of the drug. However, for at least some schizophrenia patients, knowing *how* the treatment works is just as important as knowing *whether* it does when deciding upon whether or not to take it. Cases such as this one are thus are importantly different from many others, such as the decision of whether or not to take a blood-pressure lowering medication. In the latter case, and others like it, the patient (and clinician) are most concerned with whether or not the treatment works, and very rarely, if ever, with how it does.

Given that certain groups of patients find mechanistic descriptions of treatment options an important factor to consider in their decision making process, we might wonder whether such patients *should* be given information about a treatment's supposed mechanism of action. Evidence suggests that people can be more easily persuaded to accept a given psychological claim and more likely to judge a psychological explanation to be a good explanation when the explanation is accompanied by neuroscientific evidence, even though these explanations contain reasoning errors (McCabe and Castel, 2008; Weisberg, Keil, Goodstein, Rawson, and Gray, 2008). Giving patients, particularly patients with mental illnesses, mechanistic information about a treatment option might result in the patients putting undue epistemic weight on such information in the course of their decision-making process, simply due to the mechanistic nature of the information. This could be particularly problematic when we consider the limitations in our current mechanistic knowledge. Much of our mechanistic knowledge is incomplete and chimeric in nature—cobbled together from studies in a wide range of laboratories and in different species—and thus it is unclear whether our general mechanism sketches will apply to a particular individual. If patients are given mechanistic information about their treatments, they thus might put an unwarranted amount of weight on this knowledge in the decision process. We propose that examination of this case shows that philosophers of medicine, as well as clinicians, ought to be concerned about the proper role of mechanistic reasoning in the medical decision-making process.

Mechanisms, Drug Safety and Varied Evidence

Juergen Landes

While Randomised Controlled Trials (RCTs) reign supreme in the Evidence Based Medicine (EBM) paradigm for the determination of efficacy of drugs, the role of "lower level" evidence – such as mechanistic evidence – is significantly enlarged for safety assessments. An ever growing number of philosophers of science argue also for a more prominent role of "lower level" evidence for efficacy considerations.

There are a number of reasons for the more prominent role of "lower level" evidence for drug safety within the EBM paradigm. Side effects occurring in one of ten thousand patients may be deemed unacceptable, see Food and Drug Administration (2009). Since RCTs have – virtually every – much fewer than 10:000 patients in the treatment arm(!), RCTs are much too

small to pick up rare but severe side effects. Furthermore, it is unethical to ask patients to take part in a study to establish the harmfulness of a drug.

To assess the (un-)safety of drugs it is hence imperative to take into account every bit information which could possibly offer insights into possible side effects. This presents the challenge to amalgamate statistical data and mechanistic evidence to form a view informed by all the evidence on the (un-)safety of a drug.

Frequentist statistics has no principled means to amalgamate evidence from trials and mechanistic evidence. The Bayesian approach to scientific hypothesis confirmation, e.g., Bovens and Hartmann (2003), does offer such a mean and it has been adapted for the assessment of adverse drug reactions in Landes et al. (2017).

Not only does this framework allow for a principled approach to evidence amalgamation, it also allows an analysis of epistemological value of different kinds of evidence. Philosophers have long pondered the epistemological thesis that ceteris paribus, varied evidence speaking in favor of a hypothesis confirms it more strongly than less varied evidence, Horwich (1982); Earman (1992); Wayne (1995); Steel (1996); Myrvold (1996); Fitelson (1996); Bovens and Hartmann (2003); Claveau (2013).

In this talk, I shall adapt the Bovens & Hartmann model of scientific inference to explicate the notion of varied evidence and the Variety of Evidence Thesis (VET). I will proof that all these adapted models satisfy my explication of the VET. Within these models, mechanistic evidence does hence not only play a more prominent role it also provides an epistemological value beyond its direct contribution to hypothesis confirmation. Furthermore, I shall show that bodies of evidence consisting of only one kind of evidence can only confirm a hypothesis up to a point, to go beyond the threshold evidence of other kinds is required. Interestingly, these models also pronounce on whether a diverse or a non-diverse body of evidence consisting of disconfirmatory items of evidence is more confirmatory. I shall reveal the answer to this question and argue that the models pronounce correctly. I shall argue that this strengthens the case for the VET. Given the overwhelming intuitive support in its favor and the string of negative results in Fitelson (1996); Bovens and Hartmann (2003); Claveau (2013) such good news are long overdue.

Mechanisms, disease-entities and –omics

Maël Lemoine and Brendan Clarke

In this paper, we argue that there exist clinically important models of cause and effect that are neither mechanistic nor based on evidence of difference-making (such as that arising from clinical trials). We characterise two such modes of reasoning, which we dub the **disease-entity**

model and the **–omics model**. We then contrast these with the use of mechanisms as models of disease.

We argue that the disease-entity model relies on the notion of a *fascicle*, which is an abstract account of a disease process. These fascicles are used to connect knowledge about disease aetiology to knowledge of both disease symptoms and treatments in a *narrative* way. Here, the key principle is a high degree of abstraction. This allows the model to participate in many clinically relevant inferences, such as differential diagnosis, in the absence of detailed causal knowledge. These differences in grounding (or evidence-base) provide us with our starting point from which to contrast the disease-entity model with mechanistic models of disease.

We next propose a description of the *–omic signature* model. –omic models rely on the contrast between lists of biomarkers in normal, diseased and treated states, and are intimately linked with recent developments in both automated laboratory processes and the conceptual and technological foundations of big data. We note, as a headline, that these –omics models do not depend on any reference to either causality or disease entities. Thus we give an overview of the way that these non-causal models participate in medical inferences, again by contrast with mechanistic models of disease.

In doing this, we have three aims. First, we think that the models are of philosophical interest in their own right, meaning that a positive characterisation of them is valuable for future work. Second, we want to contribute to research that follows from the Russo-Williamson thesis (and, by extension, the EBM+ project), particularly by suggesting that non-mechanistic conceptual models may play a vital role in mechanism construction. Finally, as a contribution to research on mechanisms in medicine, we would like to engage in what Mebius (2014) calls "line-drawing". We think that the current literature has paid insufficient attention to the functional boundaries of mechanisms, particularly in circumstances where actual causal knowledge is very scanty. Here, both disease-entities and –omics show how medical inferences can be sustained when mechanisms will not do the job. Not all inferences in medicine require a mechanism.

Probabilistic Causal Inference from Heterogeneous Evidence
Roland Poellinger (joint work with Juergen Landes and Barbara Osimani)

Current methods for the purpose of causal inference aim to deliver a categorical assessment about the presence of a causal relationship between events or variables. This is at odds with the great amount of epistemic and ethical uncertainty surrounding most applied sciences. In particular, for the sake of the precautionary principle, this uncertainty should not be dismissed but rather explicitly accounted for in detecting, preventing and managing, e.g., environmental or health hazards (Kreibel et al. 2001, Raffensperger, and Tickner 1999). The rationale for the attenuation of the requirement of scientific proof and certainty about the causal link is ultimately one of minimizing expected loss by anticipating risk detection and prevention (Osimani & Russo 2016; Osimani, Russo, Williamson, 2011).

We here present a framework for causal assessment which allows the incorporation of heterogeneous pieces of evidence via a probabilistic judgement about the causal link between candidate causes and effects (Landes, Osimani, Poellinger, 2017). The framework comes in the form of a Bayesian network whose nodes represent epistemic variables related to causal associations. In particular, our system i) identifies possible indicators of causality on the basis of the methodological and philosophical literature on causality, evidence, and causal inference; ii) embeds them in a topological framework of probabilistic dependencies and independencies grounded in assumptions regarding their reciprocal epistemic interconnections; iii) weakly orders some of these probabilistic dependencies as a function of their inferential strength with respect to the confirmation of causal hypotheses. This system has been developed for the purpose of drug safety assessment, but it can be easily applied to other domains with relatively few adjustments.

Our framework accommodates a number of intuitions already expressed in the literature concerning the EBM vs. pluralist debate on causal inference, evidence hierarchies, causal holism, relevance (external validity), and reliability (see for instance, Howick 2011, Clarke et al. 2014, Cartwright 2011, Teira 2011).

In this talk we will discuss how mechanistic knowledge together with statistical properties, information about difference-making, and evidence of the temporal structure can jointly support a causal hypothesis in extrapolating from study to target.

Toward a Taxonomy of Evidence of Mechanism: a Systematic Approach.

Elena Rocca

Understanding the unobservable is one of the biggest challenges in science, and the investigation of causal mechanisms in biological systems is part of it. The mechanism itself, indeed, cannot be observed directly. One has to analyze its 'signals': the effects it provokes in one or more systems. Every discussion about causal mechanisms, therefore, is intimately connected to the topic of evidence. In the field of medicine, the debate about the value of mechanistic understanding has been gaining attention. To my knowledge, however, the question of what is meant by 'evidence of mechanism' in medicine, and whether and how it can be ranked in quality, has not been systematically approached.

I propose that the challenge can be faced in two ways, which I will call respectively *fundamental* and *pragmatic*. A fundamental approach would start from a philosophical perspective, by finding an agreement about the best definition of mechanism in medicine. The agreed definition would necessarily entail some basic assumptions, which would need to be made explicit and acknowledged by all the actors at play (scientists, philosophers, clinicians, regulators, decision makers). After agreeing about the fundamental ontology, scholars could then go ahead to the epistemological discussion. These preliminary steps would minimize the risk of miscommunication, unrecognized paralogisms, and ineffective discussions. This is based on the view that a common theoretical framework generates similar empirical requirements and interpretations. Such an understanding of the scientific enterprise was started by Galilei, and held by 20th century key figures such as Einstein, Reichenbach and Kuhn.

The pragmatic approach, on the other hand, would start from an analysis of scientific patterns of discovery that successfully improved mechanistic understanding, and the comparative evaluation of the function played by different types of evidence. This should not be understood as an alternative, rather as a parallel effort to the fundamental approach. The pragmatic tactic has the advantage of the feasibility: an analysis of the scientific literature, even if extended, might require considerably less resources than the pursue of an universal theoretical agreement. However, it has the disadvantage of the external validity. The strategy would potentially identify a ranking of evidence in terms of significance, but this would be valid for some discoveries in a particular field of medicine. This might provide useful indications, however extending considerations to other cases and other fields would require caution and careful theoretical considerations.

In this paper, I use a pragmatic approach in order to review some cases of mechanistic discovery in the last three decades. I focus on the field of pharmacology, in particular harm detection (pharmacovigilance). This is because, while mechanistic understanding has a minor role when studying the efficacy of drugs, it plays an important function when establishing their safety. First, I identify successful cases of advances in understanding of causal mechanism of harm. Then, I analyze the history of discoveries and systematically classify the type of evidence that provided (i) initial signal of harm, and (ii) major advances. Finally, I analytically compare different cases. Preliminary results will be shown.

Manageable Mechanisms in Medicine: Augmenting the Russo-Williamson Thesis

Sarah Wieten

The current philosophical discussion about causes and mechanisms is diverse and robust. This literature includes debates over causal/non-causal explanations, the possibility of a causality without necessity, and mechanisms as instances of invariance (or at least constant conjunctions) (See Reutlinge; forthcoming; Mumford and Anjum, 2011, Cartwright 2009 for a sample). However, the debate over mechanisms in the philosophy of medicine is unique. Because the history of medicine is riddled with cases of the terrible (and often deadly) consequences of poor mechanistic reasoning[1], clinical medicine is currently characterized by an aversion to mechanistic approaches. In contrast, paradigms such as Evidence-Based Medicine (EBM) rely on large-scale studies of the efficacy of interventions, without the additional concern with mechanisms which lie behind this efficacy.

In this paper, I argue that in order for mechanisms to be given the important role for intervention selection that the Russo-Williamson thesis, (RWT)[2] affords them, additional theoretical work is needed. The RWT is roughly, the claim that evidence from both mechanisms and probabilistic relations is needed to warrant causal claims. This work is needed to narrow the scope of mechanisms in response to concerns raised in EBM manuals and by RWT critics, who contend that mechanisms may be so large when possible confounders are included that it is impossible to articulate them in a way that is useful for intervention decision-making. If all confounders are not accounted for, we will draw wrong, often very wrong, conclusions. In short: the worry is that what it takes to make a usable mechanistic claim that allows us to draw roughly reliable conclusions is just too demanding. Addressing the "size" of the mechanisms in question will help to bring RWT in line with concerns expressed by critics, but disagreements remain.

I develop this argument in three sections. First, I discuss the points of tension between medicine's historical distrust of mechanisms and those who insist that mechanisms have an essential role to play in medicine, such as Russo and Williamson of the RWT. I then argue that this tension stems from the assumed impossibility of understanding mechanisms in light of their immense scope. Third, I consider a possible objection to my modification of the RWT.

[1] These often cited historical cases include bloodletting for any number of ailments, anti-arrhythmic drugs for myocardial infarction, and putting babies to sleep on their stomach in order to avoid SIDS. (Howick 2011)

[2] The RWT is not at all the only thesis of its kind, with significant philosophical precursors, but given its uptake in current debate, I focus on it here.